

Harmonic analysis of cosmic microwave background data II: from ring-sets to the sky

Anthony D. Challinor,^{1*} Daniel J. Mortlock,^{1,2} Floor van Leeuwen,²
Anthony N. Lasenby,¹ Michael P. Hobson,¹ Mark A. J. Ashdown^{1,2}
and George P. Efstathiou²

¹*Astrophysics Group, Cavendish Laboratory, Madingley Road, Cambridge CB3 0HE, U.K.*

²*Institute of Astronomy, Madingley Road, Cambridge CB3 0HA, U.K.*

Accepted by MNRAS: 11 December 2001

ABSTRACT

Despite the fact that the physics of the cosmic microwave background anisotropies is most naturally expressed in Fourier space, pixelised maps are almost always used in the analysis and simulation of microwave data. A complementary approach is investigated here, in which maps are used only in the visualisation of the data, and the temperature anisotropies and polarization are only ever expressed in terms of their spherical multipoles. This approach has a number of advantages: there is no information loss (assuming a band-limited observation); deconvolution of asymmetric beam profiles and the temporal response of the instrument are naturally included; correlated noise can easily be taken into account, removing the need for additional ‘destriping’; polarization is also analysed in the same framework; and reliable estimates of the spherical multipoles of the sky and their errors are obtained directly for subsequent component separation and power spectrum estimation. The formalism required to analyse experiments which survey the full sky by scanning on circles is derived here, with particular emphasis on the *Planck* mission. A number of analytical results are obtained in the limit of simple scanning strategies. Although there are non-trivial computational obstacles to be overcome before the techniques described here can be implemented at high resolution, if these can be overcome the method should allow for a more robust return from the next generation of full-sky microwave background experiments.

Key words: cosmic microwave background – methods: analytical: – methods: numerical.

1 INTRODUCTION

The cosmic microwave background (CMB) represents the single most powerful probe of the early universe available to modern astronomy. The pioneering results of the *Cosmic Background Explorer* (COBE; Smoot et al. 1992) are being extended by the current generation of ground-based and balloon-borne experiments (e.g. Scott et al. 1996; Tanaka et al. 1996; Netterfield et al. 1997; de Oliveira-Costa et al. 1998; Coble et al. 1999; de Bernardis et al. 2000; Hanany et al. 2000; Wilson et al. 2000; Padin et al. 2001), but it is the upcoming satellite experiments that are set to revolutionise the field. By the end of 2002 the *Microwave Anisotropy Probe*[†] (MAP) will have observed the entire sky with a resolution

of $\geq 0^\circ.21$ and *Planck*[‡], scheduled for launch in 2007, should have the capacity to produce high sensitivity, all-sky maps to a resolution of $\sim 0^\circ.1$ in four of its ten frequency channels. However, it is not sufficient merely to obtain such extraordinary measurements; careful analysis of the data-sets is critical if these missions are to fulfill their promised scientific goals.

Most present techniques for CMB data analysis employ pixelised maps of the sky. Both Bond et al. (1999) and Borrill (1999) describe methods to create maximum-likelihood maps from time-ordered data and telescope pointing information, and these techniques have been successfully applied in the analysis of the BOOMERanG (de Bernardis et al. 2000; Netterfield et al. 2001) and MAXIMA (Hanany et al. 2000; Lee et al. 2001) observations. From the maps and their

* E-mail: a.d.challinor@mrao.cam.ac.uk

† <http://map.gsfc.nasa.gov/>

‡ <http://astro.estec.esa.nl/Planck/>

uncertainties, the power spectrum of the microwave sky can be estimated in a number of different ways (Bond, Jaffe & Knox 1998; Borrill 1999; Oh, Spergel & Hinshaw 1999; Wandelt, Hivon & Górski 2001), again employing maximum-likelihood techniques. These methods have a number of useful features: a map represents vast data compression relative to the time-ordered data (e.g. Borrill 1999), but is still a sufficient statistic for cosmological parameter estimation; there are usually negligible pixel-pixel noise correlations in the beam-smoothed map; maps provide important ‘reality checks’, and can be inspected by eye; and unobserved or contaminated regions of the sky can be simply removed from subsequent analysis by excluding the associated pixels (e.g. Oh et al. 1999). From a more practical point of view, the conventional map-making method is ‘tried and tested’ now. However, the use of pixelised maps during the data analysis also involves some compromises: the real microwave sky does not consist of a number of regions of uniform temperature (the prior hypothesis used to make a map), and so information is lost in map-making; subsequent steps in the analysis pipeline, such as component separation, cannot easily be performed efficiently in real space (Hobson et al. 1998); there is no single optimal or obvious choice of pixelisation scheme; and deconvolution of the (asymmetric) beam profile and temporal response of instrument, and removal of scan-synchronous instrument effects require awkward additional processing during map-making (Delabrouille 1998; Revenu et al. 2000; Wu et al. 2001). Another potential problem that has received only limited attention thus far (e.g. Tegmark & de Oliveira-Costa 2001) is how to create maps from polarization data (i.e. a map for each of the Stokes parameters, or their gradient and curl components) – there is currently no robust algorithm for the treatment of several polarization-specific systematic effects that plague CMB polarimetry experiments.

It is with these points in mind that a complementary, harmonic method of data analysis is presented here. For full-sky surveys the data analysis process would proceed from the time-ordered data directly to the set of spherical multipole moments, which would take the place of a real-space map. For band-limited observations (effectively ensured by the finite experimental beam), this harmonic reconstruction of the sky involves essentially no loss of information. The separation of the various astrophysical components of the microwave sky could then be performed quite naturally in multipole space (e.g. Hobson et al. 1998), and, if necessary, the Galactic plane could be removed while retaining orthonormality of the basis set (Górski 1994; Mortlock, Challinor & Hobson 2001; see also Section 5.2). Power spectrum estimation could then proceed naturally from this point; for instance the efficient method presented by Oh et al. (1999) could be used, although some generalisation would be required to remove the dependence on forward and inverse fast spherical transforms, which Oh et al. use to apply the inverse noise matrix efficiently in the space of multipoles and to reduce memory requirements. The multipole moments of the polarized components could be obtained using an almost identical formalism; obviously it would incur an additional overhead for each component to be reconstructed, but no further conceptual development would be required.

We have endeavoured to present the harmonic method as an end-to-end solution for the processing of time-ordered

data into power spectra. For such an analysis pipeline the benefits of clear error propagation from timelines to power spectra afforded by our harmonic methods are maximised. However, we do not pretend that all aspects of CMB analysis are best performed in the harmonic domain; instead we view harmonic methods as being complementary to standard map-based techniques. An obvious example where map-based processing is clearly desirable is the analysis of localised foregrounds (or CMB features). In addition, the inversion from observations over only a fraction of the sky to the spherical multipoles rapidly becomes singular as the resolution of the observations is increased. Although it may be possible to regularise the inversion in such cases (e.g. with prior information on the power spectrum), or to adopt a basis set adapted to the observed patch, the processing we describe in this paper is best considered in the context of full-sky observations. Even then, some aspects of the processing, such as cutting out contaminated regions close to the galactic plane for power spectrum estimation, require rather cumbersome operations if we insist on working directly with the full sky spherical multipoles rather than a map synthesised from them.

The principles of harmonic data analysis as described above are quite general, and could be applied to any full-sky CMB observations. However, it is particularly suited to experiments with a circular scanning strategy, such as the *Planck* mission. The data are obtained by a combination of rotations (i.e. the motion of the detector across the sky) and convolutions (of the beam with the true microwave sky, and this subsequent signal with the temporal response of the instrument), and these operations are most simply expressed in the harmonic basis (Delabrouille, Górski & Hivon 1998a; Challinor et al. 2000; Wandelt & Górski 2001). Hence we propose that the time-ordered data be transformed to the Fourier basis at the earliest opportunity: van Leeuwen et al. (2001) describe how to construct point source catalogues and calibrated ring-sets simultaneously, in which the data around each ring is represented in Fourier space. Some aspects of this process are necessarily mission-specific, and, for the sake of generality, it is assumed here that it is possible to obtain calibrated ring-sets (and their errors) in this form. The focus of this paper is on the subsequent ‘map’-making: combining the ring-sets to obtain an estimate of the spherical multipoles for the temperature anisotropies and polarization. In this respect, we build on earlier work by Delabrouille et al. (1998a), who analysed the statistics of the power spectrum of the Fourier modes on a single ring, and the relation with the power spectrum of the underlying temperature field on the sphere. More recently, Wandelt & Hansen (2001) have proposed the ‘ring-torus’ method for estimating the temperature power spectrum from the two-dimensional Fourier transform of data obtained on a set of rings with centres equally spaced on a small circle. While the general philosophy of the ‘ring-torus’ method coincides with that advocated here, the methodology presented here is somewhat more general. By only making use of one-dimensional Fourier data, we are able to deal with arbitrary ring-sets and to include instrument effects such as the inevitable variations in the scanning properties of the telescope (van Leeuwen et al. 2001). Furthermore, since we do not proceed directly to the power spectrum, we can make full use of Fourier-based component separation algorithms (e.g. Hobson et al. 1998) to

deal carefully with foreground contamination. The cost of maintaining this generality is the increased computational requirements of our methods over those that exploit special symmetries of the survey.

The general formalism is described in Section 2, which proceeds from a model of the instrument to the maximum-likelihood solution for the spherical multipoles of the temperature anisotropy and polarization, and their associated errors. The structure of the error covariance matrix, including polarization, is studied for some simple scan strategies in Section 3. Making use of some results given in Appendix B, we are able to make contact with several well-known analytic results for idealised experiments, previously obtained from arguments at the level of the map. In Section 4 we discuss a number of important issues that arise during the map-making phase, including correcting for scan-synchronous instrument effects, the overall calibration of the instrument to external standards, and the treatment of low frequency noise. The latter discussion includes a novel method for dealing with uncertainties in the low frequency noise power spectrum or insensitivity of the experiment to the monopole. Section 5 reviews the subsequent processing of the frequency maps, including component separation and power spectrum estimation, within the context of the harmonic data model. Finally, we conclude in Section 6 by reviewing the relative merits of the harmonic method, and suggest directions for future development.

2 THE HARMONIC DATA MODEL

The basis of the harmonic data model is the relationship between the one-dimensional Fourier representation of the ring data and the spherical multipole coefficients of the underlying sky. In the absence of instrument noise and reconstruction errors in the estimation of the Fourier ring data, the relation between the data and the sky is determined by the scanning strategy, the point-spread function or beam of the telescope, and the impulse response of the instrument in the time domain. In this section, we set up a detailed, but general, model of the instrument and scanning strategy which defines the contribution to the data coming from the sky. We then include instrument noise and reconstruction errors to give the maximum-likelihood solution for the spherical multipoles of the sky, and the covariance of their errors.

2.1 The microwave sky

We describe the microwave sky in terms of spherical multipoles on some fixed basis. The brightness in total intensity $I(\mathbf{e}; \nu)$ from the sky when observed along direction \mathbf{e} at frequency ν is represented as

$$I(\mathbf{e}; \nu) = \sum_{lm} a_{(lm)}^I(\nu) Y_{(lm)}(\mathbf{e}), \quad (1)$$

where the sum is over integers $l \geq 0$, and $|m| \leq l$. The (partial) polarization of the sky is described by Stokes (brightness) parameters $Q(\mathbf{e}; \nu)$, $U(\mathbf{e}; \nu)$, and $V(\mathbf{e}; \nu)$. We define the Q and U Stokes parameters on a basis which forms a right-handed triad with the incoming radiation direction

– \mathbf{e} : In a spherical polar coordinate system, we use the basis $\{\boldsymbol{\sigma}_\theta, -\boldsymbol{\sigma}_\phi\}$ with $\boldsymbol{\sigma}_r = \mathbf{e}$. The V brightness is a scalar function, so can be expanded in multipoles as

$$V(\mathbf{e}; \nu) = \sum_{lm} a_{(lm)}^V(\nu) Y_{(lm)}(\mathbf{e}). \quad (2)$$

It is convenient to combine the Q and U brightnesses into a symmetric, trace-free second-rank tensor:

$$\begin{aligned} \mathcal{P}^{ab}(\mathbf{e}; \nu) &= \frac{1}{2} [Q(\mathbf{e}; \nu)(\boldsymbol{\sigma}_\theta \otimes \boldsymbol{\sigma}_\theta - \boldsymbol{\sigma}_\phi \otimes \boldsymbol{\sigma}_\phi) \\ &\quad - U(\mathbf{e}; \nu)(\boldsymbol{\sigma}_\theta \otimes \boldsymbol{\sigma}_\phi + \boldsymbol{\sigma}_\phi \otimes \boldsymbol{\sigma}_\theta)], \end{aligned} \quad (3)$$

which can then be expanded in the transverse, trace-free tensor harmonics (Kamionkowski, Kosowsky & Stebbins 1997) as

$$\mathcal{P}_{ab}(\mathbf{e}; \nu) = \sum_{lm} [a_{(lm)}^G(\nu) Y_{(lm)ab}^G(\mathbf{e}) + a_{(lm)}^C(\nu) Y_{(lm)ab}^C(\mathbf{e})]. \quad (4)$$

Here, the sum is over $l \geq 2$, and $|m| \leq l$. The superscripts G (for gradient, often called electric) and C (for curl, often called magnetic) refer to the two types of transverse, trace-free harmonics. All multipoles satisfy $a_{(lm)}^{P*}(\nu) = (-1)^m a_{(l-m)}^P(\nu)$ since the brightnesses are real and we have adopted the Condon-Shortley phase for the spherical harmonics.

In this paper we assume that the $a_{(lm)}^P(\nu)$ contain all astrophysical components. In particular, we include the contribution of unresolved extra-Galactic radio sources (point sources). The point source catalogue is an important deliverable product of the *Planck* mission; an algorithm for simultaneously constructing the bright point source catalogue and calibrating the instrument from data ordered by phase on rings is described by van Leeuwen et al. (2001). Here, we assume that the contribution of those bright point sources identifiable on individual rings has not been removed from the data, although we leave open the question of whether this assumption is optimal in the context of the harmonic data model. The steps described in this paper for reconstructing the multipoles of the sky from the Fourier modes of the ring data are largely independent of whether the data includes the bright point source contribution or not. Of course, if the bright point sources are removed from the data prior to reconstructing the sky, our solution excludes the contribution from those identified point sources. One practical problem with leaving bright point sources in the ring-sets is that the sky must be represented by a larger number of multipoles to avoid biasing the estimates of the lower multipoles. Furthermore, it is then essential that the point sources are removed during the component separation phase, before any cosmological analysis of the maps (e.g. Vielva et al. 2001). Removing the brightest point sources at the level of the ring-sets leads to a modest increase in the complexity of the pipeline leading from time-ordered data to calibrated ring-sets. In particular, although a given source will only give rise to a prominent feature on a small subset of rings, the source's contribution must be removed consistently from every ring to ensure that the remaining data on each ring is derived from the same underlying sky. Point sources that are too faint to be identifiable on any single ring may still be detectable statistically when the rings are combined into the best-fitting set of multipoles (or map). Statistical detection of faint point sources is easily imple-

mented during the component separation phase (Hobson et al. 1999), and does not require the same pre-processing steps as removal of bright point sources (Vielva et al. 2001).

2.2 Ring data

A CMB mission with a circular scanning strategy can be described in terms of N_d sets of N_r rings on the sky, each set referring to a given detector (see Fig. 1). For all detectors, each ring is covered N_s times by spinning the instrument about the axis of the ring. In the case of the *Planck* mission, the satellite will rotate $N_s \simeq 60$ times about its axis before moving on to a new pointing. The r th ring is specified by its axis pointing, (θ_r, ϕ_r) , which coincides with the average pointing of the spin axis during the N_s revolutions. The axis pointing of a ring is the same for all detectors on the instrument. Note that the concept of a ring only makes sense if any precession of the spin axis (due to initial misalignment of the spin axis with the principal directions of the inertia tensor, or external torques) can be engineered to be insignificant compared to the smallest beam width. The opening angle of the r th ring for the d th detector is denoted by α_{rd} . It is defined as the average (over the N_s revolutions) angle between the spin axis and the nominal main beam direction on the sky of the appropriate horn in the focal plane. Not only does α_{rd} depend on the detector, but it also differs between rings for a given detector. The dependence on the detector is determined by the focal plane geometry. For low frequency, polarization-sensitive instruments (e.g. HEMT receivers) there will typically be two detectors (measuring nearly orthogonal polarization states) on a given horn. In such an arrangement, the ring opening angle will be the same for the two detectors, but will differ between horns. Similar comments apply to the current state of the art designs for high frequency bolometer instruments, where two orthogonal polarization sensitive bolometers are placed on a single horn. The dependence of α_{rd} within a given detector's ring-set arises from slow drifts of the spin axis relative to the instrument optics as consumables are depleted during the mission (thus changing the inertia tensor of the instrument).

The position around a given ring is specified by ψ , and is measured so that $\psi = 0$ is the most southerly intersection of the ring and the great circle containing the z -direction and the spin axis [see Fig. 1; for $\theta_r = 0$ (π) we take $\psi = 0$ to lie in the x - z plane with positive (negative) x]. The times at which different detectors pass through $\psi = 0$ is dependent on the focal plane geometry. For each horn in the focal plane we define a constant reference configuration where the nominal main beam direction is aligned with the σ_z direction of the Cartesian frame used to define the spherical multipoles for the sky, and the z -axis of the instrument reference system (van Leeuwen et al. 2001) lies in the plane normal to σ_y , with a negative projection onto σ_x . (The instrument reference system, which is fixed relative to the instrument optics, can be chosen so that its z -axis almost coincides with the nutation-averaged spin axis. Variations in the inertia tensor during the mission prevent a constant alignment of the spin and z -axes.) The instrument configuration when the d th detector takes data at angle ψ on the r th ring is obtained by rotating the instrument from the appropriate horn reference configuration. The appropriate rotation is given

by the composition[§] $D(\phi_r, \theta_r, \psi)D(0, \alpha_{rd}, \kappa_{rd})$. The rotation $D(0, \alpha_{rd}, \kappa_{rd})$ accounts for focal plane rotation (which arises from misalignment of the spin axis and the z -axis of the instrument) by: (i) first rotating the spin axis into the x - z plane of the sky coordinate system while leaving the main beam direction along the z -direction on the sky; and (ii) taking the spin axis onto the z -direction on the sky by rotating in the plane containing the spin axis and the main beam direction. The angle κ_{rd} thus measures the angle between two planes, both containing the main beam direction, with one including the spin axis and the other including the z -axis of the instrument reference system. The rotation $D(\phi_r, \theta_r, \psi)$ takes the spin axis onto the axis of the r th ring, and the main beam to the position ψ around the ring. Note that the horn reference configuration is defined by the instrument reference system rather than the spin axis. This ensures that the beam patterns in the given horn reference configuration are independent of variations in the inertia tensor.

The contribution of the (time-independent) sky signal to the time-ordered data from a given detector will be periodic in ψ , so the data can be co-added in bins of ψ with essentially no loss of useful information. Producing this phase-ordered data from the time streams is a non-trivial task, since typically the scan velocity drifts during the N_s scans of the circle: van Leeuwen et al. (2001) detail one possible scheme for reconstructing the phase-ordered data and associated errors. From the phase-ordered data $t_{rd}(\psi)$, one can estimate the Fourier coefficients in the expansion

$$t_{rd}(\psi) = \sum_n t_{(rdn)} e^{in\psi}, \quad (5)$$

and the covariance of their errors $N_{(rdn)(r'd'n')} \equiv \langle \Delta t_{(rdn)} \Delta t_{(r'd'n')}^* \rangle$. Here, angle brackets denote the expectation value and $\Delta t_{(rdn)} \equiv t_{(rdn)} - \langle t_{(rdn)} \rangle$. The $t_{(rdn)}$ are the primary data objects in the harmonic reconstruction of the microwave sky.

2.3 The instrument response

We assume that the detector responses are linear functionals of the sky in the absence of instrument noise. At any instant, the power seen by a given detector (due to the sky) is given by convolving the sky with the appropriate beam and spectral transmission. A single time-ordered data point is then obtained by convolving this power with the detector's temporal response.

2.3.1 Beam patterns

It is convenient to define the beam profile when the instrument is in the appropriate horn reference configuration, so that the beams can be assumed not to vary through the mission. A method for describing arbitrary, polarized beam patterns in terms of multipoles $b_{(lm)}^P$ was

[§] Our convention for the Euler angles α , β and γ are such that the rotation $D(\alpha, \beta, \gamma)$ actively rotates by γ about σ_z , followed by β about σ_y , and finally by α about σ_x again. All rotations are right-handed. See Brink & Satchler (1993), whose conventions we follow, for a discussion of the several alternatives that appear in the literature.




Figure 1. Geometry of the r th ring for the d th detector. The average direction of the spin axis over the N_s revolutions that make up the data on a single ring has polar angle θ_r and azimuth ϕ_r , and is the same for all detectors. The ring opening angle α_{rd} depends on the detector and the ring (due to variation in the inertia tensor during the mission). For a given ring and detector, ψ measures the angle around the ring. The point $\psi = 0$ is defined as the most southerly intersection of the ring and the plane containing the ring axis and the z -direction. Note that $\psi = 0$ is attained at different times for different detectors – the phase offsets being determined by the focal plane geometry.

introduced by Challinor et al. (2000). The beam pattern is first described in terms of effective Stokes parameters $\{\tilde{I}_d(\mathbf{e}; \nu), \tilde{Q}_d(\mathbf{e}; \nu), \tilde{U}_d(\mathbf{e}; \nu), \tilde{V}_d(\mathbf{e}; \nu)\}$, which are defined on the same $\{\boldsymbol{\sigma}_\theta, -\boldsymbol{\sigma}_\phi\}$ basis as the sky. These effective Stokes parameters can be expressed in terms of the far-field radiation pattern of the detector (Challinor et al. 2000). The $\tilde{Q}_d(\mathbf{e}; \nu)$ and $\tilde{U}_d(\mathbf{e}; \nu)$ can be combined into a linear polarization tensor for the beam, $\mathcal{B}_d^{ab}(\mathbf{e}; \nu)$, as in equation (3). The resultant power per unit frequency interval, $dW_d/d\nu$, can then be written in basis-independent form as the integral

$$\frac{d}{d\nu} W_d = A_{\text{eff},d} \int [I(\mathbf{e}; \nu) \tilde{I}_d(\mathbf{e}; \nu) - V(\mathbf{e}; \nu) \tilde{V}_d(\mathbf{e}; \nu) + 2\mathcal{P}_{ab}(\mathbf{e}; \nu) \mathcal{B}_d^{ab}(\mathbf{e}; \nu)] d\Omega, \quad (6)$$

where $A_{\text{eff},d}$ is the effective area (assumed independent of frequency). The directivity $\tilde{I}_d(\mathbf{e}; \nu)$ is a scalar field on the sphere, so can be expanded in terms of multipoles $b_{d(lm)}^I(\nu)$, as in equation (3). We impose the constraint $b_{d(00)}^I(\nu) = 1/\sqrt{4\pi}$ so that $\tilde{I}_d(\mathbf{e}; \nu)$ integrates to unity over the sphere. In a similar manner, $\tilde{V}_d(\mathbf{e}; \nu)$ can be represented by multipoles $b_{d(lm)}^V(\nu)$. Finally, $\mathcal{B}_d^{ab}(\mathbf{e}; \nu)$ can be expanded in the transverse, trace-free tensor harmonics with multipoles $b_{d(lm)}^G(\nu)$ and $b_{d(lm)}^C(\nu)$.

To obtain the power per unit frequency interval incident on the spectral filter chain for the d th detector, when the nominal main beam is at angle ψ on the r th ring, we must rotate the beam pattern by $D(\phi_r, \theta_r, \psi) D(0, \alpha_{rd}, \kappa_{rd})$

before performing the convolution in equation (6). This rotation is easily performed in multipole space to give the result (Challinor et al. 2000)

$$\frac{d}{d\nu} W_{rd}(\psi) = A_{\text{eff},d} \sum_{Plmm'm''} [a_{(lm)}^{P*}(\nu) \tilde{b}_{d(lm')}^P(\nu) \times D_{mm''}^l(\phi_r, \theta_r, \psi) D_{m''m'}^l(0, \alpha_{rd}, \kappa_{rd})], \quad (7)$$

where, for later convenience, we have defined

$$\tilde{b}_{d(lm)}^P(\nu) \equiv \delta_I^P b_{d(lm)}^I(\nu) - \delta_V^P b_{d(lm)}^V(\nu) + 2\delta_G^P b_{d(lm)}^G(\nu) + 2\delta_C^P b_{d(lm)}^C(\nu), \quad (8)$$

with $b_{d(lm)}^G(\nu) = b_{d(lm)}^C(\nu) = 0$ for $l < 2$. The $D_{mm'}^l(\alpha, \beta, \gamma)$ appearing in equation (7) are Wigner's D -matrices; our conventions follow Brink & Satchler (1993). The dependence of $D_{mm'}^l(\alpha, \beta, \gamma)$ on α and γ goes as

$$D_{mm'}^l(\alpha, \beta, \gamma) = e^{-im\alpha} d_{mm'}^l(\beta) e^{-im'\gamma}, \quad (9)$$

where the $d_{mm'}^l(\beta)$ are the reduced rotation matrices.

2.3.2 Spectral filtering

The power available at the detection element (e.g. bolometers for high frequency instruments) is given by integrating $dW_{rd}(\psi)/d\nu$ against the spectral transmission $v_d(\nu)$ of the filter chain. We shall assume that the filters on all detectors in a given frequency band are identical, and have transmission properties that are accurately known from calibrations on the ground. Map-making from multi-frequency data is usually performed on a band-by-band basis (although it is essential to use all the frequency bands to calibrate the pointing, e.g. van Leeuwen et al. 2001). Given that Fourier data from only detectors in the same frequency band are combined in such an analysis, we need only consider a single band, and henceforth assume that $v_d(\nu)$ has no dependence on the detector, i.e. $v_d(\nu) = v(\nu)$.

On integrating equation (7) against $v(\nu)$, the dependence on the sky enters through the integral $\int a_{(lm)}^{P*}(\nu) \tilde{b}_{d(lm')}^P(\nu) v(\nu) d\nu$. For the subsequent analysis, it is important that this integral can be factored into a part describing the sky carrying only l and m indices, and a part describing the beam which has only l and m' indices. Here we assume that the approximate factorisation

$$\int a_{(lm)}^{P*}(\nu) \tilde{b}_{d(lm')}^P(\nu) v(\nu) d\nu \approx \tilde{b}_{d(lm')}^P(\nu_0) \times \int a_{(lm)}^{P*}(\nu) v(\nu) d\nu, \quad (10)$$

where ν_0 is the central frequency in the band, is accurate to better than the noise level, so that we can describe the sky by the frequency integrated multipoles

$$\bar{a}_{(lm)}^P \equiv \int a_{(lm)}^P(\nu) v(\nu) d\nu. \quad (11)$$

The largest source of error resulting from this factorisation is expected to be due to the variation of the width of the main beam across the band. In Appendix A we compute the systematic error on the recovered power spectrum of the CMB temperature anisotropies resulting from this variation, assuming an axisymmetric, diffraction-limited Gaussian beam, uniform sky coverage, and no foreground contaminants. For

the *Planck* High Frequency Instrument (HFI), the bias may be non-negligible in the 100 GHz channel between $l = 1000$ and $l = 1500$, in which case refinements to equation (10) will be required. Here, we adopt the factorisation in equation (10), in which case the power available at the detection element, $W_{rd}(\psi)$, can be written as

$$W_{rd}(\psi) = A_{\text{eff},d} \sum_{Plmm'} [\bar{a}_{(lm)}^{P*} D_{mm'}^l(\phi_r, \theta_r, \psi) \times \mathcal{B}_{rd(lm')}^P]. \quad (12)$$

For later convenience, we have introduced the notation

$$\mathcal{B}_{rd(lm')}^P \equiv \sum_{m''} D_{m'm''}^l(0, \alpha_{rd}, \kappa_{rd}) \tilde{b}_{d(lm'')}^P(\nu_0) \quad (13)$$

for the multipoles of the d th beam after rotation by $D(0, \alpha_{rd}, \kappa_{rd})$. Note that the r -dependence of $\mathcal{B}_{rd(lm')}^P$ derives only from the (slow) variation of opening angle α_{rd} and focal plane rotation κ_{rd} through the mission.

2.3.3 Temporal response

The available power, $W_{rd}[\psi(t')]$, is convolved with the temporal response function of the detector, $h_d(t, t')$, to generate a single time-ordered datum at time t . We decompose $h_d(t, t')$ into a stationary part, $h_d(t - t')$, and a stochastic part, $\Delta h_d(t, t')$, which we assume to have stationary statistical properties. The stochastic part, which derives from e.g. random fluctuations in amplifier gains, will be considered as part of the instrument noise in Section 2.4. Here, we concentrate on the stationary impulse response function $h_d(t)$. Typically, this will contain several instrument artifacts, such as the effect of finite sampling and non-zero detector time constants, which are convolved together to form the total impulse response. For example, if the detector time constant is τ_d , and samples are taken by integrating over time intervals Δ_d and are assigned to the midpoint of the interval, then $h_d(t) = h_{\tau_d, d}(t) \star h_{\Delta_d, d}(t)$, where the detector impulse response

$$h_{\tau_d, d}(t) = G_d \tau_d^{-1} e^{-t/\tau_d} \Theta(t/\tau_d), \quad (14)$$

and the sampling impulse response

$$h_{\Delta_d, d}(t) = \Delta_d^{-1} \Theta(1 + 2t/\Delta_d) \Theta(1 - 2t/\Delta_d). \quad (15)$$

Here, $\Theta(x)$ is the Heaviside unit step function, and G_d is the average gain over the mission, which includes both amplifier gain and the quantum efficiency of the detectors.

One of the advantages of describing the data on rings in the Fourier domain is that the convolution with the temporal response of the instrument reduces to a simple product (Delabrouille et al. 1998a). Denoting the contribution of the signal to the phase-ordered data by $s_{rd}(\psi)$, we have

$$s_{rd}(\psi) = \int_{-\infty}^{\infty} W_{rd}(\psi + \omega_r t') h_d(-t') dt', \quad (16)$$

where $\omega_r = d\psi/dt$ is the (average) spin rate on the r th ring. We have assumed that the support of $h_d(t)$ is sufficiently compact that the only contribution to $s_{rd}(\psi)$ is from the available power W_{rd} on the same ring. We have further assumed that any variation of the spin rate about the average is negligible for the purpose of performing the convolution,

so that the convolution of the available power with the impulse response is periodic in ψ . Inserting equation (12) into equation (16), and expanding $D_{mm'}^l(\phi_r, \theta_r, \psi)$ as in equation (9), we find

$$s_{rd}(\psi) = A_{\text{eff},d} \sum_{Plmm'} [\bar{a}_{(lm)}^{P*} D_{mm'}^l(\phi_r, \theta_r, \psi) \times H_d^*(m' \omega_r) \mathcal{B}_{rd(lm')}^P], \quad (17)$$

where $H_d(\omega_r)$ is the Fourier transform of the impulse response:

$$H_d(\omega_r) = \int_{-\infty}^{\infty} h_d(t) e^{-i\omega_r t} dt. \quad (18)$$

For the impulse responses in equations (14) and (15), we have

$$H_d(\omega_r) = \frac{G_d \text{sinc}(\omega_r \Delta_d/2)}{1 + i\omega_r \tau_d}. \quad (19)$$

Finally, we can extract from equation (17) the contribution of the signal to the n th Fourier coefficient of the phase-ordered data. Complex conjugating and using equation (9), we find

$$s_{(rdn)} = A_{\text{eff},d} \sum_{Plm} [\bar{a}_{(lm)}^P d_{mn}^l(\theta_r) e^{im\phi_r} H_d(n\omega_r) \mathcal{B}_{rd(lm)}^{P*}], \quad (20)$$

where the sum is over $l \geq |n|$ and $|m| \leq l$. Note that the effect of the impulse response can also be described in terms of an effective (ring-dependent) beam with multipoles

$$b_{rd(lm)}^{P, \text{eff}}(\nu) = \sum_{m'm''} [D_{mm'}^l(-\kappa_{rd}, -\alpha_{rd}, 0) H_d^*(m' \omega_r) \times D_{m'm''}^l(0, \alpha_{rd}, \kappa_{rd}) b_{d(lm'')}^P(\nu)]. \quad (21)$$

The effective beam will be generally not be axisymmetric as a consequence of the smearing along the scan direction induced by the impulse response of the instrument. The concept of an effective beam becomes particularly significant if the beam has to be calibrated from the observations of e.g. bright point source transits, as described by van Leeuwen et al. (2001). In such cases, it is the effective beam that can be reconstructed most directly. A potential difficulty that would need to be addressed in such an approach is the (slow) variation of effective beam with ring due to variations in the angles α_{rd} and κ_{rd} . In practice, it may be simpler to deconvolve the temporal response of the instrument prior to construction of the phase-ordered data. In this case, $H_d(n\omega_r)$ can be omitted from equation (20).

2.4 Gaussian random noise and reconstruction errors

The Fourier modes, $t_{(rdn)}$, that we construct from the phase-ordered data differ from the $s_{(rdn)}$ of equation (20) because of a number of sources of error. For example, stochasticity in the instrument temporal response, such as that due to thermal noise in the detectors and amplifiers, has not been included in equation (20). Neither have we included the effects of photon (shot) noise, or any detector offsets. Furthermore, we have replaced the spin axis pointing by its average over the N_s spin periods, ignoring the (albeit

small) effects of the nutation of the instrument. These random errors tend to be suppressed in the phase-ordered data due to the averaging over N_s periods, except for fluctuations at temporal frequencies synchronous with the spin frequency. Fluctuations below the spin frequency remain in the phase-ordered data predominantly as an offset on the ring. In the usual map-making paradigm (Bond et al. 1999; Borrill 1999), low frequency ($1/f$) noise must be carefully accounted for; failure to do so leads to highly correlated errors in the reconstructed map which typically appear as stripes (e.g. Delabrouille 1998; Revenu et al. 2000). It is difficult to accommodate low frequency noise power in map-making from the time-ordered data since its inclusion spoils the sparse nature of the time-time noise correlation matrix. Approximate methods to circumvent this problem for experiments that scan on rings already exist (Delabrouille 1998; Revenu et al. 2000); these ‘destriping’ methods will be discussed further in Section 4.1.

There are also a number of potential systematic errors which contaminate the $t_{(rdn)}$, and, if left unaccounted for, will lead to bias in the reconstructed multipoles of the sky. Keeping the numerous systematics under control is essential to maintaining the integrity of the final data products. A significant merit of the harmonic data model is that the model of the instrument, equation (20), is sufficiently complete to include a number of systematic effects which are not so naturally included in the standard map-making techniques. For example, the harmonic model can seamlessly accommodate any prior knowledge we have of asymmetries in the beam profiles, cross-polar contamination for the polarized detectors, sidelobe features leading to ‘straylight’ entering the instrument, and detector time constants and data sampling rate. For some systematics there will be incomplete (or even no) prior knowledge of essential parameters, in which case these parameters must be estimated from the data themselves during the reconstruction of the sky. The harmonic model also appears well-suited to handling such systematics iteratively; see Section 4.2.

To include the effect of Gaussian random noise with a known (or estimated) power spectrum in the reconstruction of the sky from the Fourier data on rings, $t_{(rdn)}$, we require the covariance matrix $N_{(rdn)(r'd'n')}$. This matrix was calculated by Delabrouille et al. (1998a) for a single detector and ring ($d = d'$ and $r = r'$). Here we extend their result to include the effects of a finite sampling rate (see also Janssen et al. 1996), and to allow for $r \neq r'$. Assuming instrument noise dominates the random errors $\Delta t_{(rdn)}$, there will be negligible correlation between the errors on different detectors, so we need only consider $d = d'$.

We assume that the noise contribution, $n_d(t)$, to the time stream of the d th detector is a real, stationary random process with zero mean, and power spectrum $\mathcal{N}_d(\omega)$. The power spectrum is the Fourier transform of the noise auto-correlation function $C_d(\tau) = \langle n_d(t + \tau)n_d(t) \rangle$:

$$\mathcal{N}_d(\omega) = \int_{-\infty}^{\infty} C_d(\tau) e^{-i\omega\tau} d\tau. \quad (22)$$

The propagation of the noise $n_d(t)$ to the errors $\Delta t_{(rdn)}$ depends on the exact procedure for estimating the Fourier modes $t_{(rdn)}$ from the time-ordered data (van Leeuwen et al. 2001). However, to a first approximation, the errors $\Delta t_{(rdn)}$ are obtained by convolving $n_d(t)$ with the impulse response

of the sampler, $h_{\Delta_d,d}(t)$, mapping the portion of this function covering the observation interval for the r th ring onto the angle ψ around the ring, and finally extracting the Fourier coefficients of the resultant signal:

$$\Delta t_{(rdn)} = \frac{\omega_r}{2\pi} \int_{t_{rd}^-}^{t_{rd}^+} h_{\Delta_d,d} \star n_d(t) \exp[-in\omega_r(t - t_{rd}^-)] dt. \quad (23)$$

Here, t_{rd}^- is the time when the d th detector is first pointing to $\psi = 0$ on the r th ring, t_{rd}^+ is the end of the observation interval for that ring: $t_{rd}^+ = t_{rd}^- + 2\pi N_s/\omega_r$, and we have ignored any variation of the spin rate during the N_s revolutions. It is now straightforward to write the covariance of the errors in terms of the noise power spectrum. For the case of scanning missions such as *Planck*, the number of revolutions $N_s \gg 1$ is sufficiently large that negligible correlations remain between different Fourier modes. In this case we find

$$\begin{aligned} N_{(rdn)(r'd'n')} &= \delta_{dd'} \delta_{nn'} \frac{1}{2\pi} \int_{-\infty}^{\infty} \{\mathcal{N}_d(\omega) \text{sinc}^2(\omega\Delta_d/2) \\ &\quad \times \text{sinc}^2[(\omega/\omega_r - n)\pi N_s] \\ &\quad \times \exp[i2\pi N_s(r - r')\omega/\omega_r]\} d\omega, \end{aligned} \quad (24)$$

where we have further ignored any variation of the average spin rate, ω_r , between rings. Provided that $\mathcal{N}_d(\omega) \text{sinc}^2(\omega\Delta_d/2)$ varies slowly compared to the rest of the integrand in equation (24) around $\omega = n\omega_r$, we can replace it by its value there. The remaining integral vanishes unless $r = r'$, so that for slowly varying power spectra the covariance matrix is fully diagonal:

$$\begin{aligned} N_{(rdn)(r'd'n')} &= \delta_{rr'} \delta_{dd'} \delta_{nn'} \\ &\quad \times \frac{\omega_r}{2\pi N_s} \mathcal{N}_d(n\omega_r) \text{sinc}^2(n\omega_r\Delta_d/2). \end{aligned} \quad (25)$$

In the presence of a significant low frequency noise component (such as $1/f$ noise in the electronics) joining onto white noise at a knee frequency ω_{knee} (e.g. Delabrouille 1998), the assumption of a slowly varying noise power spectrum may not hold for small $|n|$. Typically the spin rate will be chosen so that $\omega_{\text{knee}} \ll \omega_r$, in which case equation (25) will be valid for $n \neq 0$. However, for $n = 0$ we can expect correlations between rings unless $\omega_r/N_s \ll \omega_{\text{knee}} \ll \omega_r$. (For the *Planck* HFI, the nominal $\omega_{\text{knee}} < 0.06 \text{ rad s}^{-1}$, and $\omega_r \approx 0.10 \text{ rad s}^{-1}$, so these conditions are likely to be satisfied.) Even in the presence of such correlations, the noise covariance matrix is still block diagonal in the harmonic representation, in contrast to the time-time covariance matrix which has significant off-diagonal terms due to the extended support of the noise auto-correlation function, $C_d(\tau)$.

2.5 Maximum-likelihood solution

Our model for the ring Fourier data $t_{(rdn)}$ is now $t_{(rdn)} = s_{(rdn)} + \Delta t_{(rdn)}$, where the signal $s_{(rdn)}$ is modelled by equation (20). It is convenient to write this relation in the form

$$t_{(rdn)} = \sum_{Plm} A_{(rdn)(Plm)} \bar{a}_{(lm)}^P + \Delta t_{(rdn)}, \quad (26)$$

where the coupling matrix is

$$A_{(rdn)(Plm)} \equiv A_{\text{eff},d}^l m_n(\theta_r) e^{im\phi_r} H_d(n\omega_r) \mathcal{B}_{rd(ln)}^{P*}. \quad (27)$$

We define $A_{(rdn)(Plm)}$ to be zero for $l < |n|$, since the $t_{(rdn)}$ only depend on the $\bar{a}_{(lm)}^P$ with $l \geq |n|$. This structure can

be exploited to derive an unbiased (although not minimal variance) estimator for the $\bar{a}_{(lm)}^P$, which can be computed in $O(l_{\max}^4)$ operations. Here, l_{\max} is the maximum l that we retain during the inversion. (The experimental beam limits l_{\max} even if the underlying sky is not band-limited.) The estimate is derived by working down from the maximum Fourier mode, n_{\max} , and performing a regularised inversion of a subset of the data to solve for those $\bar{a}_{(lm)}^P$ on which the data subset depends but which have not been determined at a previous step. Full details of this estimator will be given in a future paper (Mortlock et al. in preparation).

In this paper we shall just give the formal maximum-likelihood solution to equation (26). We assume that we are attempting to solve only for the $\bar{a}_{(lm)}^P$ which form the components of a vector \mathbf{a} . In practice, it may also be desirable to attempt to solve for several instrument parameters which are unknown from earlier calibrations, but which represent significant systematic effects. Such parameters are easily included in the maximum-likelihood formalism, but unless they influence the data linearly their inclusion would prevent us from being able to locate the maximum-likelihood multipoles analytically. Given the size of the problem, having to locate the maximum of the likelihood by a numerical search in model space is clearly undesirable. The treatment of systematics with unknown parameters is probably best performed iteratively (e.g. Delabrouille, Gispert & Puget 1998b; see also Section 4.2).

Assuming Gaussian noise, the maximum-likelihood estimate, $\hat{\mathbf{a}}$, for the true sky, \mathbf{a} , is

$$\hat{\mathbf{a}} = (\mathbf{A}^\dagger \mathbf{N}^{-1} \mathbf{A})^{-1} \mathbf{A}^\dagger \mathbf{N}^{-1} \mathbf{t}, \quad (28)$$

where \mathbf{t} is the vector of Fourier modes $t_{(rdn)}$, \mathbf{A} is the matrix of coupling coefficients $A_{(rdn)(r'd'n')}$, and \mathbf{N} is the noise covariance matrix with components $N_{(rdn)(r'd'n')}$. The † operation in equation (28) denotes Hermitian conjugation. The maximum-likelihood solution is the optimal, unbiased estimate of the sky. Its covariance matrix is

$$\mathbf{C} \equiv \langle (\hat{\mathbf{a}} - \mathbf{a})(\hat{\mathbf{a}} - \mathbf{a})^\dagger \rangle = (\mathbf{A}^\dagger \mathbf{N}^{-1} \mathbf{A})^{-1}, \quad (29)$$

where the average is over noise realisations. A brute force evaluation of $\hat{\mathbf{a}}$ requires $O(l_{\max}^6)$ operations and $O(l_{\max}^4)$ storage which is impossible at the moment for high resolution experiments. Conjugate gradient techniques (e.g. Press et al. 1992) would reduce the operations count to $O(N_i l_{\max}^4)$, where N_i is the number of iterations required, by removing the matrix multiplies and inversion. Note that the block-diagonal structure of \mathbf{N} (equation 24) reduces the cost of computing \mathbf{N}^{-1} to $O(l_{\max}^4)$ operations, or fewer if correlations between rings are confined to the $n = 0$ modes. Keeping N_i small requires a careful choice of preconditioner, i.e. a Hermitian, positive-definite matrix which is both a good approximation to the inverse covariance matrix \mathbf{C}^{-1} and is easy to invert. We demonstrate in Section 3 that for simple scan strategies such matrices can easily be found (see also Oh et al. 1999). For the case where all rings are at similar latitude, an approximation to \mathbf{C}^{-1} can be computed in $O(l_{\max}^4)$ operations and requires only $O(l_{\max}^3)$ storage. This preconditioner is block-diagonal and can be inverted in $O(l_{\max}^4)$ operations. Note, however, that the conjugate gradient evaluation of equation (29) still requires $O(l_{\max}^4)$ storage for the non-sparse matrix \mathbf{A} . The application of the conjugate gradient method, as well other iterative techniques, to equa-

tion (29) will be explored numerically in Mortlock et al. (in preparation).

In writing down the maximum-likelihood solution we have assumed that $\mathbf{A}^\dagger \mathbf{N}^{-1} \mathbf{A}$ is invertible, so that the solution is unique. For some scan strategies and instrument geometries, $\mathbf{A}^\dagger \mathbf{N}^{-1} \mathbf{A}$ may become numerically singular due to incomplete coverage of the sky. With incomplete coverage, and l_{\max} sufficiently large, there may exist sets of multipoles which produce a temperature or polarization field which is localised to machine precision in the regions of the sky that are not covered (Mortlock et al. 2001). In such cases, solving for the multipoles must proceed via singular value techniques, and we obtain no constraint on the part of \mathbf{a} which lies in the null space of the coupling matrix \mathbf{A} . If map-making were performed in pixel space, similar problems would arise when attempting to estimate the multipoles from the pixelised map. As an alternative to singular-value techniques, we could further regularise the inversion by Wiener filtering. With a Gaussian prior for the signal \mathbf{a} , with covariance $\mathbf{S} \equiv \langle \mathbf{a} \mathbf{a}^\dagger \rangle$, the maximum of the posterior probability gives the Wiener-filtered reconstruction

$$\hat{\mathbf{a}} = (\mathbf{S}^{-1} + \mathbf{A}^\dagger \mathbf{N}^{-1} \mathbf{A})^{-1} \mathbf{A}^\dagger \mathbf{N}^{-1} \mathbf{t}. \quad (30)$$

In this manner, the inversion of the singular matrix $\mathbf{A}^\dagger \mathbf{N}^{-1} \mathbf{A}$ is regularised by the inverse signal covariance matrix \mathbf{S}^{-1} . In practice, the microwave sky is only poorly approximated as a Gaussian random process, particularly in those frequency channels where the CMB is not dominant. In addition, we will only have limited knowledge of the signal covariance matrix \mathbf{S} either from other observations, or a preliminary, approximate analysis of the data. However, Wiener filtering has proved to be a useful technique in component separation where similar objections to the use of Gaussian priors hold (see e.g. Hobson et al. 1998 and references therein), and the same should be true for harmonic ‘map’-making. Note that Wiener filtering produces a biased estimate of the sky; other linear filters have been proposed to circumvent this problem, but generally produce noisier, though unbiased, maps (e.g. Tegmark & Efstathiou 1996). Equation (30) can also be solved with conjugate gradient techniques; in this case the preconditioner should include some simplified form of the signal covariance in addition to an approximation to $\mathbf{A}^\dagger \mathbf{N}^{-1} \mathbf{A}$.

In our discussion so far we have assumed that full beam deconvolution is performed during the ‘map’-making stage. In practice, deconvolving the beam completely may be undesirable for several reasons. Firstly, the matrix $\mathbf{A}^\dagger \mathbf{N}^{-1} \mathbf{A}$ is likely to be more ill-conditioned since the errors on the reconstructed multipoles must grow large as l approaches the (inverse) beam scale. (In a conjugate gradients solution, the preconditioner can partly take this effect into account.) Secondly, as discussed in Section 2.3.2 and Appendix A, variation of the width of the main beam across the spectral band of a channel is a potential source of systematic error. For an asymmetric beam, one could make the optimistic assumption beam multipoles in a given frequency channel could be written as the product of a frequency-dependent, symmetric part $\tilde{b}_{(l0)}^P(\nu)$ which is assumed to be the same for all detectors in the band, and a frequency-independent, asymmetric part which may depend on the detector:

$$\tilde{b}_{d(lm)}^P(\nu) \approx \tilde{b}_{(l0)}^P(\nu) [\tilde{b}_{d(lm)}^P(\nu_0) / \tilde{b}_{(l0)}^P(\nu_0)], \quad (31)$$

where ν_0 is the central frequency in the band. We could then solve for the multipoles of the sky convolved with an effective symmetric beam with multipoles $\tilde{b}_{(lm)}^P(\nu)$ and integrated across the frequency band with filter $v(\nu)$. By factoring the beam in this way, we can simultaneously improve the condition number of $\mathbf{A}^\dagger \mathbf{N}^{-1} \mathbf{A}$ and remove most of the bias due to variation of the main beam width with frequency.

3 SIMPLE SCAN STRATEGIES

In this section we investigate the structure of the inverse covariance matrix, \mathbf{C}^{-1} , of the recovered multipoles for two simple scanning strategies. The first is a random pointing of the spin axis for each ring, which ensures uniform coverage of the entire sky. The second is constant latitude scanning, where all rings have the same polar angle, $\theta_r = \theta$. Constant latitude scanning is a useful approximation to the scan strategy of the *Planck* satellite, where the spin axis stays close to the plane of the ecliptic. By adopting coarse restrictions on the behaviour of the instrument, we are able to demonstrate that, under these restrictions, the harmonic estimate of the sky, equation (28), is statistically equivalent to that obtained via conventional map-making and spherical analysis. By relaxing some restrictions on the behaviour of the instrument, we obtain some useful extensions of the known results.

Throughout this section we impose a number of limits on the instrument response: (i) The spin axis position is always aligned with the z -axis of the instrument reference system (see Section 2.2), so that the ring opening angles, α_{rd} , are independent of r , and the focal plane rotations $\kappa_{rd} = 0$; (ii) There is no variation in average spin rate between rings, so we can drop the subscript from ω_r ; (iii) The band width of the spectral filters is sufficiently narrow that

$$\int a_{(lm)}^{P*}(\nu) \tilde{b}_{(lm')}^P(\nu) v(\nu) d\nu = a_{(lm)}^{P*}(\nu_0) \tilde{b}_{(lm')}^P(\nu_0) \times \int v(\nu) d\nu, \quad (32)$$

where ν_0 is the central frequency of the band; (iv) Bolometer time constants, τ_d , and the sampling period, Δ_d , are such that $\tau_d \omega_r l_{\max} \ll 1$ and $\Delta_d \omega_r l_{\max} \ll 1$, in which case the effect of the instrument response can be approximated by simply a gain at zero lag; (v) The noise power spectrum is white so that

$$N_{(rdn)(r'd'n')} = \delta_{rr'} \delta_{dd'} \delta_{nn'} \sigma_d^2, \quad (33)$$

where σ_d^2 is a constant for each detector, which can be related to the detector sensitivity, s_d .[¶] With the restrictions above,

$$s_d = \left[A_{\text{eff},d} G_d \bar{T} \frac{\partial B(\nu_0, T)}{\partial T} \right]_{T=\bar{T}} \int v(\nu) d\nu \Bigg]^{-1}$$

[¶] The detector sensitivity is defined such that for an integration time Δt with a single pointing, the excess signal to noise is unity when the instrument is illuminated with an unpolarized black-body sky with a temperature excess $Ts/\sqrt{\Delta t}$ over the average CMB temperature, \bar{T} . If the units of the gain G_d are e.g. VW^{-1} , so that $s_{rd}(\psi)$ is a readout Voltage, σ_d will also have the dimensions of a Voltage, and s_d will have units of e.g. $(\mu\text{K}/\text{K})\text{s}^{1/2}$.

$$\times \sqrt{\frac{2\pi N_s}{\omega_r}} \sigma_d, \quad (34)$$

where $B(\nu, T)$ is the Planck function and the average CMB temperature is \bar{T} . Note that a consequence of restriction (i) is that $\mathcal{B}_{rd(l'n)}^P$ is independent of ring, $\mathcal{B}_{rd(l'n)}^P = \mathcal{B}_{d(l'n)}^P$. If we define dimensionless multipoles by

$$\tilde{a}_{(lm)}^P \equiv \left[\bar{T} \frac{\partial B(\nu_0, T)}{\partial T} \right]_{T=\bar{T}}^{-1} a_{(lm)}^P(\nu_0), \quad (35)$$

the inverse of the dimensionless error covariance matrix, $\tilde{\mathbf{C}}^{-1}$, has components

$$\tilde{C}_{(Plm)(P'l'm')}^{-1} = \frac{2\pi N_s}{\omega} \sum_{rdn} s_d^{-2} \{ D_{mn}^l(\phi_r, \theta_r, 0) \mathcal{B}_{d(l'n)}^P \times [D_{m'n}^{l'}(\phi_r, \theta_r, 0) \mathcal{B}_{d(l'n)}^{P'}]^* \}, \quad (36)$$

where the sum is over detectors, rings, and $|n| \leq \min(l, l')$. We now proceed to analyse $\tilde{\mathbf{C}}^{-1}$ for randomly-positioned rings and constant latitude scanning.

3.1 Randomly-positioned rings

We consider the limit of equation (36) as the number of rings $N_r \rightarrow \infty$ while keeping the survey length finite, in which case we can replace the sum over rings by an appropriate integral. In practice, for finite N_r the continuum limit should hold for scales with $l \lesssim \sqrt{N_r/(4\pi)}$. Placing the rings at random is equivalent to enforcing uniform coverage of the full sky, which will allow us to make contact with existing analytic results obtained from the map-based formalism. The sum over randomly-positioned rings in equation (36) can be replaced by $[N_r/(4\pi)] \int d\Omega_r$, where $d\Omega_r \equiv d\phi_r d\cos\theta_r$. Making use of the orthogonality of the $D_{mm'}^l(\alpha, \beta, \gamma)$ over the $\text{SO}(3)$ group manifold (Brink & Satchler 1993), we find

$$\int D_{mn}^l(\phi_r, \theta_r, 0) [D_{m'n}^{l'}(\phi_r, \theta_r, 0)]^* d\Omega_r = \frac{4\pi}{2l+1} \delta_{ll'} \delta_{mm'}, \quad (37)$$

so that $\tilde{\mathbf{C}}^{-1}$ simplifies to

$$\tilde{C}_{(Plm)(P'l'm')}^{-1} = \delta_{ll'} \delta_{mm'} \frac{T_m}{2l+1} \sum_{dn} s_d^{-2} \mathcal{B}_{d(l'n)}^P \mathcal{B}_{d(l'n)}^{P'*}, \quad (38)$$

where the total mission time is $T_m = 2\pi N_s N_r / \omega$. Equation (38) can be further simplified by using equation (13) to substitute for $\mathcal{B}_{d(l'n)}^P$ and then performing the sum over n noting that $[D_{mn}^l(\alpha, \beta, \gamma)]^* = D_{nm}^l(-\gamma, -\beta, -\alpha)$. The result is

$$\tilde{C}_{(Plm)(P'l'm')}^{-1} = \delta_{ll'} \delta_{mm'} \frac{T_m}{2l+1} \times \sum_{dm''} s_d^{-2} \tilde{b}_{d(lm'')}^P(\nu_0) \tilde{b}_{d(lm'')}^{P'*}(\nu_0). \quad (39)$$

Note that there is now no dependence on the opening angles of the detectors, $\alpha_{rd} = \alpha_d$. This is a consequence of assuming white noise. If the noise had a characteristic timescale, this would combine with the spin velocity and opening angle to define a characteristic angular scale for the projected noise, and $\tilde{\mathbf{C}}$ would then depend on α_d .

To make further progress, we consider initially the case where the beams are axisymmetric, and have no cross-polar contamination. In this limit, the intensity multipoles can be written in terms of a window function, $W_{l,d}$:

$$b_{d(lm)}^I(\nu_0) = \sqrt{\frac{2l+1}{4\pi}} W_{l,d} \delta_{m0}, \quad (40)$$

and the linear polarization multipoles in terms of spin-weight 2 window functions, ${}_2W_{l,d}$, (Challinor et al. 2000; see also Ng & Liu 1999):

$$b_{d(lm)}^G(\nu_0) = -\sqrt{\frac{2l+1}{32\pi}} {}_2W_{l,d} e^{-im\rho_d} (\delta_{m2} + \delta_{m-2}), \quad (41)$$

$$b_{d(lm)}^C(\nu_0) = -i\sqrt{\frac{2l+1}{32\pi}} {}_2W_{l,d} e^{-im\rho_d} (\delta_{m2} - \delta_{m-2}). \quad (42)$$

Here, ρ_d is the angle between the polarization direction on axis and the normal to the plane containing the main beam direction and the spin axis. (The unit normal to this plane is σ_y when the detector is in its horn's reference configuration). By assumption, the circular polarization multipoles for the beam vanish. Substituting into equation (39), we find

$$\begin{aligned} \tilde{C}_{(Plm)(P'l'm')}^{-1} &= \delta_{ll'} \delta_{mm'} \delta_{PP'} \\ &\times \sum_d w_d [\delta_I^P W_{l,d}^2 + (\delta_G^P + \delta_C^P) {}_2W_{l,d}^2], \end{aligned} \quad (43)$$

where we have introduced the weight per solid angle (Knox 1995), which for uniform coverage of the sky is $w_d \equiv T_m/(4\pi s_d^2)$. Note that the covariance matrix is diagonal, and that it does not depend on the relative orientations of the polarization directions of the various detectors. This is not surprising since for randomly-positioned rings, every point on the sky is traversed by every detector in every possible orientation, and we have assumed the instrument noise to be uncorrelated between detectors.

It is straightforward to show that the covariance matrix in equation (43) is equal to that obtained by constructing optimal (minimum variance), beam-smoothed maps of the Stokes parameters for each detector, and then optimally extracting the $\tilde{a}_{(lm)}^P$ from a joint analysis of the maps. If the main beam of the d th detector lies in the p th pixel when the ring phase is ψ on the r th ring, the signal contribution to the phase-ordered data, equation (17), can be written as (see e.g. Challinor et al. 2000)

$$\begin{aligned} s_{rd}(\psi) &= A_{\text{eff},d} G_d \bar{T} \left. \frac{\partial B(\nu_0, T)}{\partial T} \right|_{T=\bar{T}} \int v(\nu) d\nu [I_{\text{eff},d}(\mathbf{e}_p) \\ &- Q_{\text{eff},d}(\mathbf{e}_p) \cos 2\eta + U_{\text{eff},d}(\mathbf{e}_p) \sin 2\eta]. \end{aligned} \quad (44)$$

under the restrictions described above. In equation (44), \mathbf{e}_p is in the direction of the p th pixel, and has polar angle θ_p and azimuth ϕ_p , and the angle η is the angle between the polarization direction (on the sky) along the beam axis and the plane containing \mathbf{e}_p and the z -axis of the fixed reference system. The beam-smoothed fields are

$$I_{\text{eff},d}(\mathbf{e}) = \sum_{lm} W_{l,d} \tilde{a}_{(lm)}^I Y_{(lm)}(\mathbf{e}) \quad (45)$$

$$\begin{aligned} \frac{1}{\sqrt{2}}(Q_{\text{eff},d} \pm iU_{\text{eff},d})(\mathbf{e}) &= \sum_{lm} [{}_2W_{l,d} (\tilde{a}_{(lm)}^G \mp i\tilde{a}_{(lm)}^C) \\ &\times \mp {}_2Y_{(lm)}(\mathbf{e})], \end{aligned} \quad (46)$$

where the ${}_2Y_{(lm)}$ are examples of the spin-weight harmonics, defined for integer s by (Goldberg et al. 1967)

$$D_{m-s}^l(\alpha, \beta, \gamma) = (-1)^s \sqrt{\frac{4\pi}{2l+1}} s Y_{(lm)}^*(\beta, \alpha) e^{is\gamma}. \quad (47)$$

For randomly-positioned rings in the limit $N_r \rightarrow \infty$, the angles ψ are uniformly covered in any given pixel. Given our assumptions, the optimal, unbiased estimate for $I_{\text{eff},d}$ in the p th pixel is given by direct averaging of the appropriate $t_{rd}(\psi)$. For $Q_{\text{eff},d}$ the $t_{rd}(\psi)$ are averaged with weight $-2\cos 2\eta$, and for $U_{\text{eff},d}$ the weight is $2\sin 2\eta$. The errors on the smoothed maps are uncorrelated between Stokes parameters, and between pixels. The errors on these maps have weights per solid angle w_d for $I_{\text{eff},d}$, and $w_d/2$ for $Q_{\text{eff},d}$ and $U_{\text{eff},d}$.

For uncorrelated noise between Stokes parameters, pixels, and detectors, the maximum-likelihood estimate of the $\tilde{a}_{(lm)}^P$ from the smoothed maps reduces to minimising the (absolute) squared residuals between the left and right-hand sides of equations (45) and (46), with each pixel for each detector entering with the appropriate weight per solid angle. In the case of uniform coverage considered in this subsection, the maximum-likelihood estimates of the $\tilde{a}_{(lm)}^P$ are equivalent to performing a spherical transform of the smoothed maps and then averaging these across the detectors, with each detector carrying statistical weight $W_{l,d} w_d$. However, it will be useful for the next subsection to note the form that the error covariance matrix takes when we relax the assumption of uniform coverage, while retaining the assumptions of uncorrelated errors between Stokes parameters, detectors and pixels. The weights per solid angle for $I_{\text{eff},d}$ are then pixel-dependent, $w_{d,p}$, and the weights for $Q_{\text{eff},d}$ and $U_{\text{eff},d}$ are $w_{d,p}/2$. The non-vanishing entries of the inverse of the (Hermitian) error covariance matrix are then:

$$\tilde{C}_{(Ilm)(Il'm')}^{-1} = \sum_{pd} W_{l,d} W_{l',d} w_{d,p} Y_{(lm)}^* Y_{(l'm')} \Delta\Omega_p, \quad (48)$$

$$\begin{aligned} \tilde{C}_{(Glm)(Gl'm')}^{-1} &= \sum_{pd} [{}_2W_{l,d} {}_2W_{l',d} w_{d,p} \frac{1}{2} ({}_2Y_{(lm)}^* {}_2Y_{(l'm')}) \\ &+ {}_2Y_{(lm)}^* {}_2Y_{(l'm')}) \Delta\Omega_p], \end{aligned} \quad (49)$$

$$\begin{aligned} \tilde{C}_{(C lm)(Cl'm')}^{-1} &= i \sum_{pd} [{}_2W_{l,d} {}_2W_{l',d} w_{d,p} \frac{1}{2} ({}_2Y_{(lm)}^* {}_2Y_{(l'm')}) \\ &- {}_2Y_{(lm)}^* {}_2Y_{(l'm')}) \Delta\Omega_p], \end{aligned} \quad (50)$$

and $\tilde{C}_{(C lm)(Cl'm')}^{-1} = \tilde{C}_{(G lm)(Gl'm')}^{-1}$. Here, $\Delta\Omega_p$ is the solid angle subtended by the p th pixel. For the case of uniform coverage, $w_{d,p} = w_d$, the orthonormality of the spin-weight harmonics reduces equations (48)–(50) to the result derived from the harmonic model, equation (43), if we ignore pixelisation effects. For this simple example, the estimates of the $\tilde{a}_{(lm)}^P$ from the harmonic model and the map-making route are both unbiased and have the same error covariance. It follows that they are statistically equivalent at the level of second-order correlations (and at all orders for Gaussian noise).

3.1.1 Effect of beam asymmetry

We now consider the effect of a known beam asymmetry on the covariance matrices obtained via the harmonic and map-based routes. For simplicity, we shall only consider unpolarized detectors with bivariate Gaussian profiles, with eccentricity e_d . Consider initially a detector in its horn's refer-

ence configuration, with the major axis of the iso-directivity ellipse aligned with σ_x . For beam widths $\sigma_d \ll 1$, we have

$$\tilde{I}_d(e; \nu_0) \approx \frac{\exp\{-\theta^2[\cos^2 \phi + \sin^2 \phi/(1 - e_d^2)]/(2\sigma_d^2)\}}{2\pi\sigma_d^2\sqrt{1 - e_d^2}}, \quad (51)$$

which has non-zero multipoles, in the limit of large l ,

$$b_{d(lm)}^I = \sqrt{\frac{2l+1}{4\pi}} I_{m/2}(l^2\sigma_d^2 e_d^2/4) e^{-l^2\sigma_d^2(1-e_d^2/2)/2}, \quad (52)$$

for m even, where $I_m(x)$ are modified Bessel functions. Note that beam asymmetry is only significant for $l \gg 1/(\sigma_d e_d)$. The multipoles for a detector whose iso-directivity ellipse has major axis at angle γ_d to σ_x pick up an additional phase $\exp(-im\gamma_d)$.

The expression for $\tilde{C}_{(l'm')^{-1}}^{(l'm)}$ (equation 39) involves the sum $\sum_{|n| \leq l} |b_{d(l'n)}^I(\nu_0)|^2$. For large l , this is easily evaluated by substituting the integral representation of the modified Bessel functions (or, equivalently, azimuthally averaging the absolute square of the flat-space Fourier transform of the directivity). The result is

$$\sum_{|n| \leq l} |b_{d(l'n)}^I(\nu_0)|^2 = \frac{2l+1}{4\pi} I_0(l^2\sigma_d^2 e_d^2/2) e^{-l^2\sigma_d^2(1-e_d^2/2)/2}, \quad (53)$$

so that

$$\tilde{C}_{(l'm')^{-1}}^{(l'm)} = \delta_{ll'} \delta_{mm'} \sum_d \left[w_d I_0(l^2\sigma_d^2 e_d^2/2) \times e^{-l^2\sigma_d^2(1-e_d^2/2)/2} \right]. \quad (54)$$

If a known beam asymmetry is to be accounted for with pixel-based map-making techniques, it is necessary to include the asymmetry in the real-space pointing matrix. A simpler procedure suggests itself for the case of randomly positioned rings, where we can exploit the fact that every pixel is sampled with each detector in every orientation. Averaging the data in each pixel from a given detector gives an unbiased estimate of the pixelised sky smoothed with an effective window function. This window function follows from azimuthally averaging the beam, so that

$$W_{\text{eff},l,d} \equiv \sqrt{\frac{4\pi}{2l+1}} b_{d(l0)}^I = I_0(l^2\sigma_d^2 e_d^2/4) e^{-l^2\sigma_d^2(1-e_d^2/2)/2}. \quad (55)$$

The analysis of the smoothed maps proceeds as in Section 3.1, so that the inverse covariance matrix evaluates to

$$\tilde{C}_{(l'm')^{-1}}^{(l'm)} = \delta_{ll'} \delta_{mm'} \sum_d \left[w_d I_0^2(l^2\sigma_d^2 e_d^2/4) \times e^{-l^2\sigma_d^2(1-e_d^2/2)} \right]. \quad (56)$$

Clearly there is some information loss on averaging the data in the manner described. For a single detector, the ratios of the diagonal elements of the covariance matrices obtained from equations (54) and (56) are

$$\frac{I_0^2(l^2\sigma_d^2 e_d^2/4)}{I_0(l^2\sigma_d^2 e_d^2/2)} = 1 - \frac{1}{32} (l\sigma_d e_d)^4 + \dots \quad (57)$$

This information loss is only significant on scales below that of the beam asymmetry ($\approx \sigma_d e_d$). For more general scanning, averaging data in a pixel will necessarily produce a

biased estimate of the sky in the presence of a beam asymmetry. This will be most acute for scanning strategies where a large number of pixels are sampled with the detectors in only a narrow range of orientations. Such strategies include constant latitude scanning, to which we now turn.

3.2 Constant latitude rings

To analyse constant latitude scanning, we consider equation (36) under the restriction $\theta_r = \theta$. Following the discussion at the start of Section 3.1, we consider the continuum limit. For constant latitude scans we replace the sum over rings by $[N_r/(2\pi)] \int d\phi_r$. Expanding the D -matrices as in equation (9) and performing the integral over ϕ_r , the expression for $\tilde{\mathbf{C}}^{-1}$ reduces to

$$\tilde{C}_{(Plm)(P'l'm')}^{-1} = 4\pi \delta_{mm'} \sum_{dn} w_d d_{mn}^l(\theta) d_{m'n}^{l'}(\theta) \mathcal{B}_{d(l'n)}^P \mathcal{B}_{d(l'n')}^{P'*}, \quad (58)$$

where the sum is over detectors and $|n| \leq \min(l, l')$. We obtain uncorrelated errors between modes with different m due to the azimuthal symmetry of the sky coverage (Oh et al. 1999). Since $\tilde{\mathbf{C}}^{-1}$ is block-diagonal, it can be inverted in $O(l_{\text{max}}^4)$ operations. In practice, variations in the instrument parameters through the mission will spoil this exact block-diagonal structure, as will any precession in the latitude θ . However, approximating $\tilde{\mathbf{C}}^{-1}$ by its diagonal blocks may still provide an adequate preconditioner for e.g. a conjugate gradient reconstruction of the sky.

In equation (58), $w_d \equiv T_m/(4\pi s_d^2)$ refers to the weight per solid angle for uniform coverage of the whole sky, which differs from the pixel-dependent $w_{d,p}$ due to the variation in time spent per solid angle. By differentiating the geometric relation

$$\cos \theta_p = \cos \alpha_d \cos \theta - \sin \alpha_d \sin \theta \cos \psi, \quad (59)$$

which relates the polar angle θ_p of the pixel containing the main beam to the ring phase ψ , it is straightforward to show that

$$w_{d,p} = \frac{2}{\pi} \frac{\Theta(B_{d,p})}{\sqrt{B_{d,p}}} w_d, \quad (60)$$

where $\Theta(x)$ is the Heaviside unit step function, and we have introduced the quantity

$$B_{d,p} \equiv 1 - \cos^2 \alpha_d - \cos^2 \theta_p - \cos^2 \theta + 2 \cos \alpha_d \cos \theta_p \cos \theta \quad (61)$$

for later convenience.

To proceed further, we make the assumption that the detector sensitivities and ring opening angles are all equal, so that $w_d = w$, $\alpha_d = \alpha$, and $B_{d,p} = B_p$. It follows that the pixel-dependent weight per solid angle is also independent of detector: $w_{d,p} = w_p$. The sum over detectors in equation (58) then reduces to

$$\sum_d \mathcal{B}_{d(l'n)}^{P*} \mathcal{B}_{d(l'n')}^{P'} = \sum_{m''m'''} [d_{nm''}^l(\alpha) d_{nm'''}^{l'}(\alpha) \times \sum_d \tilde{b}_{d(l'm'')}^P \tilde{b}_{d(l'm''')}^{P'*}]. \quad (62)$$

We further assume that all detectors have axisymmetric, co-polar beams which are equivalent up to rotations about the beam axis. In this case, the multipoles are given by

equations (40)–(42) with the angles ρ_d determining the relative polarization orientations. The choice of angles ρ_d has a strong impact on the covariance structure of the estimated multipoles. Designing the focal plane so that $\sum_d \exp(2i\rho_d) = 0$ ensures that the errors are uncorrelated between the total intensity and the linear polarization. The correlation between the G and C multipoles is controlled in part by the sum $\sum_d \exp(4i\rho_d)$, although demanding that this sum vanish is not a sufficient condition to ensure uncorrelated errors between G and C . However, the condition $\sum_d \exp(4i\rho_d) = 0$ does ensure uncorrelated errors between the smoothed Stokes fields^{||} $Q_{\text{eff}}(e)$ and $U_{\text{eff}}(e)$ which, under the stringent restrictions adopted in this section, can be estimated in an optimal manner by a χ^2 fitting of the data from all detectors obtained when the main beam of each lies in the given pixel (Couchot et al. 1999). The smoothed maps estimated in this manner have uncorrelated errors between pixels with weights per solid angle $N_d w_p$ for I_{eff} , and $N_d w_p/2$ for Q_{eff} and U_{eff} , where, recall, N_d is the number of detectors. Following Couchot et al. (1999), we denote arrangements of polarimeters with $\sum_d \exp(2i\rho_d) = 0$ and $\sum_d \exp(4i\rho_d) = 0$ as optimal configurations. In such a configuration, the non-vanishing components of the inverse covariance matrix evaluate to

$$\tilde{C}_{(lmI)(l'm'I)}^{-1} = \delta_{mm'} w N_d W_l W_{l'} \sqrt{(2l+1)(2l'+1)} \times \sum_n d_{mn}^l(\theta) d_{m'n}^{l'}(\theta) d_{n0}^l(\alpha) d_{n0}^{l'}(\alpha), \quad (63)$$

$$\tilde{C}_{(lmG)(l'm'G)}^{-1} = \delta_{mm'} w N_d/2 W_l/2 W_{l'} \sqrt{(2l+1)(2l'+1)} \times \sum_n \{d_{mn}^l(\theta) d_{m'n}^{l'}(\theta) \frac{1}{2} [d_{n2}^l(\alpha) d_{n2}^{l'}(\alpha) + d_{n-2}^l(\alpha) d_{n-2}^{l'}(\alpha)]\}, \quad (64)$$

$$\tilde{C}_{(lmG)(l'm'C)}^{-1} = -i \delta_{mm'} w N_d/2 W_l/2 W_{l'} \sqrt{(2l+1)(2l'+1)} \times \sum_n \{d_{mn}^l(\theta) d_{m'n}^{l'}(\theta) \frac{1}{2} [d_{n2}^l(\alpha) d_{n2}^{l'}(\alpha) - d_{n-2}^l(\alpha) d_{n-2}^{l'}(\alpha)]\}, \quad (65)$$

with $\tilde{C}_{(lmC)(l'm'C)}^{-1} = \tilde{C}_{(lmG)(l'm'G)}^{-1}$. In equations (63)–(65), the sums are over $|n| \leq \min(l, l')$.

We show in Appendix B that the sum over n on the right-hand sides of equations (63)–(65) can be reduced to matrix elements of the time spent per solid angle with the appropriate spin weight basis:

$$\delta_{mm'} \sum_n \sqrt{(2l+1)(2l'+1)} d_{mn}^l(\theta) d_{m'n}^{l'}(\theta) d_{ns}^l(\alpha) d_{ns}^{l'}(\alpha) = \int \frac{2}{\pi} \frac{\Theta(B_p)}{\sqrt{B_p}} {}_{-s}Y_{(lm)}^*(e_p) {}_{-s}Y_{(l'm')}(e_p) d\Omega_p, \quad (66)$$

for integer s . Using this result in equations (63)–(65), and recalling equation (60), the components of the inverse co-

variance matrix can be reduced to

$$\tilde{C}_{(lmI)(l'm'I)}^{-1} = W_l W_{l'} \int w_p Y_{(lm)}^* Y_{(l'm')} d\Omega_p, \quad (67)$$

$$\tilde{C}_{(lmG)(l'm'G)}^{-1} = 2 W_l/2 W_{l'} \int [w_p \frac{1}{2} (2 Y_{(lm)}^* 2 Y_{(l'm')} + 2 Y_{(lm)}^* {}_{-2}Y_{(l'm')}) d\Omega_p], \quad (68)$$

$$\tilde{C}_{(lmG)(l'm'C)}^{-1} = i 2 W_l/2 W_{l'} \int [w_p \frac{1}{2} (2 Y_{(lm)}^* 2 Y_{(l'm')} - 2 Y_{(lm)}^* {}_{-2}Y_{(l'm')}) d\Omega_p]. \quad (69)$$

Note that these results automatically take account of incomplete sky coverage, through the presence of $\Theta(B_p)$. The function $B_p < 0$ in pixels that are not sampled by the main beam during the mission, so such pixels make no contribution to the covariance matrix. As noted in Section 2.5, the presence of unsampled regions can render $\tilde{\mathbf{C}}^{-1}$ numerically singular. For the *Planck* mission the nominal ring opening angle $\alpha \approx 85^\circ$, which would leave 5° holes at the ecliptic poles for constant latitude scanning in the ecliptic. However, a proposal to precess the spin axis out of the ecliptic plane with an amplitude of 10° has the benefit of providing complete sky coverage, which would render $\tilde{\mathbf{C}}^{-1}$ invertible (although no longer block-diagonal).

To compare equations (67)–(69) with the results obtained from a pixelised map, we can make use of equations (48)–(50) which, it will be recalled, give the errors on the maximum-likelihood solution for the multipoles obtained from N_d sets of smoothed maps. Here, we have only one map (estimated using data from all detectors as described above), with weights per pixel $N_d w_p$ for I_{eff} and $N_d w_p/2$ for Q_{eff} and U_{eff} . With these conditions, it is straightforward to verify that equations (48)–(50) reduce to pixelised versions of equations (67)–(69). This observation confirms the statistical equivalence of the map-based and harmonic routes through to the multipoles of the sky for constant latitude scanning, under the strong restrictions adopted in this section.

More generally, since the maps and their multipoles contain the same information (disregarding pixelisation errors), it will always be the case that the optimal map-making and map-to-multipole algorithms are statistically equivalent to the optimal harmonic estimate of the multipoles. However, as demonstrated in Section 2, the harmonic model provides a more natural framework for the inclusion of a number of systematic effects in the instrument modelling and data analysis pipeline. The inclusion of some such effects is essential to maintain the integrity of the final data products.

4 DISCUSSION

In this section we discuss a number of important issues that arise during the map-making stage. The treatment of low frequency noise (‘destriping’) in the harmonic model is described in Section 4.1. In Section 4.2 we comment on the control of scan-synchronous instrument effects, focusing on straylight from bright sources picked up through the side-lobes of the telescope, and the modulation of the dipole in the rest-frame of the experiment. Finally, the overall calibration of the experiment to external standards is described in Section 4.3.

^{||} The smoothed fields are defined in equations (45) and (46). We have dropped the subscript d since we are assuming that the beam window functions are the same for all detectors.

4.1 Destriping

In Section 2.4 we established that for large N_s , the main impact of noise power at frequencies below the spin frequency is concentrated at the $n = 0$ Fourier modes of the phase-ordered data (see also, e.g. Delabrouille 1998). If the noise power varies sufficiently rapidly in the range $0 \leq \omega \leq \omega_r/N_s$ these offsets may be correlated between rings. Careful treatment of low frequency noise is essential to avoid undesirable long range noise correlations ('stripes') in the maps. Given the difficulty of inverting the time-time noise covariance matrix in the presence of low frequency noise, additional pre-processing steps are usually performed prior to map-making. Pre-whitening the time- or phase-ordered data (e.g. Tegmark 1997) uses a prior knowledge of the noise power spectrum, $\mathcal{N}_d(\omega)$, to represent the data in a basis where there are no noise correlations. In essence, this involves the subtraction of an optimal offset from each ring of data. Natoli et al. (2001) have successfully applied this method to the 30 GHz channel of the *Planck* Low Frequency Instrument under the assumption of a symmetric beam profile. Other destriping methods advocated for the *Planck* mission, e.g. Delabrouille (1998) and Revenu et al. (2000), also involve the subtraction of offsets from the rings, but these offsets are obtained directly from the data by exploiting the intersections of rings.

The harmonic model represents the phase-ordered data in the Fourier basis on rings which ensures that the noise covariance matrix, \mathbf{N} , is very sparse, even in the presence of significant low frequency power. This removes the need for any additional pre-whitening. A potential problem arises if the noise power spectrum is not known accurately at low frequencies, since this may compromise the quality of the sky reconstruction. The harmonic model suggests a simple solution to this problem: remove the $n = 0$ Fourier modes from the analysis. From a Bayesian viewpoint, we must now infer the $\bar{a}_{(lm)}^P$ from a knowledge of $t_{(rdn)}$ with $n > 0$. This inversion requires the probability density function (pdf) for $n_{(rdn)}$ with $n > 0$, which is obtained from the full pdf by integrating over the $n = 0$ modes. However, for Gaussian noise, the pdf factors into products of pdfs for the individual Fourier modes as a consequence of equation (24), so that integrating over the $n = 0$ modes is trivial. Note also that since the $n = 0$ modes of the data are the only ones that depend on the $l = 0$ modes of the total intensity and circular polarization, these monopole modes can no longer be reconstructed. To implement the scheme one can either reformulate the problem with the $n = 0$ data modes removed at the outset, or, equivalently, let $N_{(rd0)(r'd0)} \rightarrow \infty$, and perform the inversion in the subspace orthogonal to $\bar{a}_{(00)}^I$ and $\bar{a}_{(00)}^V$. Removing the $n = 0$ modes does not bias the solution, but it is no longer optimal given all the data, since $t_{(rd0)}$ receives a contribution from all multipoles. However, for large l_{\max} the information loss should be small for the $l > 0$ multipoles. Excluding the $n = 0$ modes provides a very robust, simple way of removing undesirable long range correlations in the maps that might arise if the low frequency noise were estimated inaccurately. Note also that this 'destriping' method handles the polarization signal seamlessly, avoiding the technicalities involved in estimating offsets from ring intersections for polarized detectors (Revenu et al. 2000). Removing the $n = 0$ and $l = 0$ modes from the analysis also neatly solves the

problem of degeneracy that arises if the experiment is insensitive to the monopole – this situation may arise for the *Planck* HFI since the favoured readout electronic system is insensitive to very low frequencies (Gaertner et al. 1997).

4.2 Control of Scan-Synchronous Instrument Effects

Instrument effects that produce signals synchronous with the rotation of the instrument will not be suppressed by co-adding the time-ordered data to form the phase-ordered data. Given complete knowledge of the instrument response, these potential sources of systematic error can be controlled by adopting a more refined model of the instrument response during the reconstruction of the sky. In this subsection we discuss the removal of two potential systematic effects within the context of the harmonic data model.

4.2.1 Sidelobe Reconstruction

An important source of systematic error is straylight from e.g. the Galaxy and the CMB dipole entering the instrument through the sidelobes of the telescope. Here, the difficulty lies not in the inclusion of the sidelobes in the reconstruction process, since the harmonic model allows for a complete description of the beams, but rather in the limited knowledge of the sidelobes that will be available from simulations and from in-flight calibrations (Burigana et al. 2001; van Leeuwen et al. 2001). Although the sidelobes can be expected to be highly polarized, the direction should be sufficiently random that it is adequate to model only the directivity $\tilde{I}_d(\mathbf{e}; \nu_0)$ in the sidelobes.

Delabrouille et al. (1998b) have proposed an iterative scheme for estimating sidelobe corrections. In their method, an estimate of the sky is obtained from a first pass of the map-making algorithm ignoring sidelobe corrections, and this is then used to compute the difference between the observed time-ordered data and the contribution of the sky coming through the main beam. This difference is attributable to instrument noise and the sidelobe signal, and can be inverted (with suitable regularisation) to estimate the values of the directivity in the sidelobe pixels. With this improved knowledge of the beam, and the original estimate of the sky, an estimate of the sidelobe contribution can be subtracted from the time-ordered data, and an improved estimate of the sky obtained from a further pass of the original map-making algorithm (again with sidelobes ignored). This process can be iterated to give consistent estimates of the sky and the directivity in the sidelobe. It should be possible to implement a similar scheme in the harmonic model, with the sidelobes parameterised by a modest number of multipoles and the sidelobe contribution estimated from the $t_{(rdn)}$, although we have not attempted this yet. One potential problem with such a scheme is that the sidelobes may contain rather sharp features due to geometrical effects, in which case a parameterisation of the sidelobes with spherical multipoles may not be ideal. The monopole $b_{(00)}^I(\nu_0)$ should not be varied during sidelobe reconstruction since its value is fixed at $1/\sqrt{4\pi}$ by definition. The iterative reconstruction of the sidelobes does not require an absolute calibration of the detector gains, and should be performed prior to the calibration procedures described in Section 4.3.

4.2.2 Dipole variation

One subtlety that we have overlooked so far, which is significant for survey missions from space, is the variation of satellites (linear) velocity during the year. The orbital velocity of the earth about the sun $\sim 30 \text{ km s}^{-1}$, which induces a yearly modulation in the multipoles seen in the satellite's rest frame. For sensitivities equivalent to the *Planck* mission (a few $\mu\text{K}/K$ in 100 arcmin^2 pixels at 100 GHz, for twelve months of observation), the most significant effect is the modulation of the intensity dipole due to the monopole with amplitude $\sim 10^{-4} K$. If we denote the multipoles in the frame of the satellite by $a_{(lm),E}^P(\nu)$, and continue to denote the multipoles on the $\{\sigma_x, \sigma_y, \sigma_z\}$ basis fixed relative to the solar system by $a_{(lm)}^P(\nu)$, we have

$$a_{(lm),E}^P(\nu) \approx a_{(lm)}^P(\nu) + \delta_I^P \delta_{l1} \frac{\nu^4}{\sqrt{4\pi}} \frac{d}{d\nu} \left[\frac{a_{(00)}^I(\nu)}{\nu^3} \right] \times \int \mathbf{v} \cdot \mathbf{e} Y_{(lm)}^*(\mathbf{e}) d\Omega, \quad (70)$$

where \mathbf{v} is the (time-dependent) velocity of the earth relative to the sun, and \mathbf{e} is a unit vector. The $a_{(lm),E}^P(\nu)$ are defined relative to a basis obtained by Lorentz boosting the $\{\sigma_x, \sigma_y, \sigma_z\}$ basis. Aberration effects on the high l multipoles may also be significant for *Planck* (Challinor & van Leeuwen, in preparation).

It is the time-dependent multipoles $\bar{a}_{(lm),E}^P$ (obtained by averaging $a_{(lm),E}^P(\nu)$ across the spectral filter) which should now appear in equation (20), relating the signal to the sky. However, we must still solve for the time-independent multipoles $\bar{a}_{(lm)}^P$. If the monopole $a_{(00)}^I(\nu)$ were available at the start of the analysis for all frequency channels, the contribution of the time-dependent part of $a_{(lm),E}^P(\nu)$ to the signal could be subtracted from the $t_{(rdn)}$ prior to solving for the multipoles. Note that this process only affects the $n = 0$ and $n = 1$ Fourier modes, $t_{(rdn)}$, but these influence all multipoles of the reconstructed sky. More realistically, the monopole $a_{(00)}^I(\nu)$ may only be poorly determined prior to the analysis, in which case it will be necessary to iterate the estimation of the multipoles, using the monopole determined in the previous iteration to correct the $t_{(rdn)}$. This process is clearly not possible if the $n = 0$ modes are ignored in the analysis, as may be necessary for the reasons discussed in Section 4.1. A radical way to circumvent the problems of poorly determined low frequency noise and dipole modulation is to remove the $n = 1$ modes from the analysis also. However, an internally reconstructed dipole appears to be essential for establishing an absolute normalisation of the experiment (see next subsection).

4.3 Gain calibration

The methods described by van Leeuwen et al. (2001) for producing internally consistent ring-sets do not provide an absolute gain calibration for each detector. By forgoing the introduction of any external flux calibrator during the reconstruction of the ring-sets, only the ratios of the gains of all detectors in a given frequency band will be known; the absolute gains will only be determined up to an overall factor, g . This factor is defined to be the ratio of the true gains, $G_{d,\text{true}}$, to those assumed in the model of the instrument, G_d

(e.g. equation 19). An external calibrator provides prior information on the sky, $\text{Pr}(\mathbf{a})$, which can be used to constrain g . For simplicity, we assume a Gaussian prior:

$$\text{Pr}(\mathbf{a}) \propto |\pi \mathbf{C}_{\text{ext}}|^{-1/2} \exp[-(\mathbf{a} - \mathbf{a}_{\text{ext}})^\dagger \mathbf{C}_{\text{ext}}^{-1} (\mathbf{a} - \mathbf{a}_{\text{ext}})/2], \quad (71)$$

where \mathbf{a}_{ext} is the most likely prior sky, and \mathbf{C}_{ext} is the associated error. The sky and normalisation can be now estimated directly from the (Fourier) ring data, \mathbf{t} , by maximising $\text{Pr}(\mathbf{t}|\mathbf{a}, g)\text{Pr}(\mathbf{a})$ (assuming a uniform prior on g). In the usual limit where the experiment determines the sky to much better precision than the calibrator, the best estimate of the sky is $\hat{\mathbf{a}}/\hat{g}$, where $\hat{\mathbf{a}}$ is the maximum-likelihood estimate for the sky given in equation (28), obtained with gains G_d , and the best estimate of the normalisation is given by

$$\hat{g} = \frac{\hat{\mathbf{a}}^\dagger \mathbf{C}_{\text{ext}}^{-1} \hat{\mathbf{a}}}{\Re(\hat{\mathbf{a}}^\dagger \mathbf{C}_{\text{ext}}^{-1} \mathbf{a}_{\text{ext}})}. \quad (72)$$

In practice, $\mathbf{C}_{\text{ext}}^{-1} \mathbf{a}_{\text{ext}}$ and $\mathbf{C}_{\text{ext}}^{-1}$ may have to be estimated from existing observations on patches of the sky. In such cases, equation (72) reduces to the solution obtained by minimising the (error-weighted) residuals between the reconstructed map, $\sum_{lm} g^{-1} \hat{a}_{(lm)}^I Y_{(lm)}(\mathbf{e})$, and the calibration map. The requirement that the calibration maps be at the same frequencies as the instrument channels can be removed by selecting a large patch, at high Galactic latitude, where the large-scale signal is dominated by the CMB dipole.

5 FURTHER ANALYSIS OF FREQUENCY MAPS

Although the main focus of this paper is the reconstruction of frequency maps of the sky in multipole space, in this section we offer some comments on the subsequent analysis of these maps into their astrophysical components, and estimated power spectra.

5.1 Component separation

The frequency maps contain contributions from a number of astrophysical components which must be separated on the basis of their assumed frequency spectra (and possibly power spectra). Unresolved radio sources require separate processing since they cannot be accurately modelled as a population with a single frequency spectrum. If bright point sources were not removed from the ring-sets prior to map-making, their contribution should be filtered from the multipoles before attempting component separation (Vielva et al. 2001). Within the timescale of *Planck*, the majority of point sources expected to be visible in the 857 GHz channel of the HFI should already have been catalogued prior to launch by the *ASTRO-F* survey^{★★}. Such catalogues will be valuable for the geometrical calibration of the ring-sets using the positions of identified point-sources (van Leeuwen et al. 2001), and also for the removal of point sources from the reconstructed frequency maps. If bright point sources have already been removed from the ring-sets prior to map-making, component separation can proceed directly from the reconstructed multipoles. However, statistical reconstruction of

★★ <http://www.ir.isas.ac.jp/ASTRO-F/>

faint point sources may still be desirable, in which case the faint point sources can be included as a generalised noise in the separation algorithm for the other astrophysical components (Hobson et al. 1999).

The full-sky, maximum-entropy component separation algorithm developed recently by Stolyarov et al. (2001) performs the separation in the spherical multipole basis, so that it can use the outputs of the multipole estimation process described here directly. Although the algorithm of Hobson et al. does not handle polarization in its current form, the extension to polarized components should be straightforward. (The Wiener separation of polarized components has been implemented successfully by Bouchet, Prunet & Sethi 1999 for small patches of sky.) A more demanding problem for high resolution, all-sky separation algorithms is the inclusion of non-isotropic noise, since then the separation cannot be performed multipole by multipole due to the correlated errors (e.g. Prunet et al. 2001).

5.2 Power spectrum estimation

Power spectrum estimation can also be performed efficiently in multipole space. The simplest, unbiased estimator which is quadratic in the estimated dimensionless multipoles \tilde{a}_{lm}^P is of the form

$$\hat{C}_l^{PP'} = \frac{1}{(2l+1)} \sum_m (\tilde{a}_{lm}^P \tilde{a}_{lm}^{P'*} - \tilde{C}_{(Plm)(P'lm)}). \quad (73)$$

The variance of this estimator is calculated in Kamionkowski et al. (1997) under the assumptions of full sky coverage and isotropic pixel noise which is uncorrelated between Stokes parameters. In this case, the quadratic estimator in equation (73) is equivalent to the maximum-likelihood estimator.

For more general noise properties and scan strategies the estimator in equation (73) is sub-optimal. Maximum-likelihood techniques have been developed which allow accurate computation of the temperature power spectrum in $O(l_{\max}^4)$ operations for scanning strategies with (near) azimuthal symmetry (Oh et al. 1999). This method cannot be directly in the context of the harmonic model since Oh et al. store the inverse (noise) covariance matrix, \mathbf{C}^{-1} , in the pixel basis only, where they assume it is diagonal. Forward and inverse fast spherical transforms are then employed to apply \mathbf{C}^{-1} efficiently in multipole space. Applying \mathbf{C}^{-1} directly in multipole space would not increase the overall operations count significantly, but would increase the memory requirements to $O(l_{\max}^4)$ in general.

One potential issue in power spectrum estimation is the treatment of regions near to the Galactic plane. Although component separation algorithms can reconstruct the CMB very well even at low Galactic latitude (Stolyarov et al. 2001), it may still be necessary to remove traces of the Galaxy by brute force. Given a pixelised map, the pixels near the plane can be excised from the analysis before estimating the multipoles and their error covariance, and subsequent power spectrum estimation. (Alternatively, power spectrum estimation can be performed directly in real space away from the Galactic plane, although the lack of a sparse signal covariance matrix in the pixel representation is problematic for high resolution data-sets.)

Within the context of the harmonic model, the Galactic plane can be dealt with by projecting the reconstructed

multipoles, \hat{a}_{lm}^P , into a subspace which is (almost) free from Galactic contamination. For simplicity, we shall only describe the procedure for a total power measurement; the extension to polarized data is straightforward. We define the Hermitian matrix \mathbf{P} to have components

$$P_{(lm)(l'm')} \equiv \int_{S^{2'}} Y_{lm}^*(\mathbf{e}) Y_{l'm'}(\mathbf{e}) d\Omega, \quad (74)$$

where the integral is over the region of the sky that we wish to retain in the analysis. In the limit $l_{\max} \rightarrow \infty$, $P_{(lm)(l'm')}$ are the matrix elements of a projection operator (i.e. $\mathbf{P}^2 = \mathbf{P}$) which projects functions into the region $S^{2'}$. For finite $l_{\max} \gg 1$, \mathbf{P} is almost a projection operator since the distribution of its eigenvalues is almost bimodal with values clustered close to zero and one (Mortlock et al. 2001). The fraction of eigenvalues close to unity is approximately the fraction of the sky retained; the remainder are nearly all zero. Performing (maximum-likelihood) power spectrum estimation with the data object $\mathbf{P}\hat{\mathbf{a}}$ would not remove the Galactic contamination since \mathbf{P} is (formally) invertible. To enforce rejection of the Galaxy, we introduce a projection operator $\tilde{\mathbf{P}}$ which is obtained from \mathbf{P} by retaining the eigenvectors but setting those eigenvalues greater than $1 - \epsilon$ to unity, and to zero otherwise. The choice of threshold, $\epsilon \ll 1$, must be determined from simulations to ensure good rejection of the Galaxy while minimising the number of modes lost from the analysis. Working with $\tilde{\mathbf{P}}\hat{\mathbf{a}}$ ensures that we have projected out those modes of $\hat{\mathbf{a}}$ which are (nearly) localised in the region external to $S^{2'}$. This procedure is similar to power spectrum estimation from the cut map, since projecting pixelised data into $S^{2'}$ in real space has the effect of greatly amplifying the noise on the multipoles estimated from the cut map along those directions in multipole space which correspond to functions nearly localised in the cut. In practice, it is likely to be more efficient to represent the projected data vector $\tilde{\mathbf{P}}\hat{\mathbf{a}}$ on a basis adapted to the region $S^{2'}$, so that we work with the object $\mathbf{U}^\dagger \hat{\mathbf{a}}$, where \mathbf{U} is the (non-square) matrix whose columns are the eigenvectors of $\tilde{\mathbf{P}}$ with unit eigenvalue. The matrix \mathbf{U} can easily be obtained from a singular value decomposition of $\tilde{\mathbf{P}}$ (Mortlock et al. 2001).

6 CONCLUSION

Most current techniques for analysing CMB data on the sphere are based on the use of pixelised maps. The basis for a complementary approach, based on spherical harmonic coefficients of the intensity and polarization, has been derived here for the case of experiments that scan the full sky in circles. Our method offers a number of advantages, most notably the refined treatment of non-ideal beam effects, the ability to handle correlated (low frequency) noise in an optimal, but efficient, manner, and the seamless way in which polarized data can be analysed alongside total intensity data. In principle, harmonic methods can be used to develop an end-to-end pipeline for the analysis of full-sky survey data, allowing the clean propagation of noise and other errors from the time-ordered data to the CMB power spectra. In practice, the methods described here are best regarded as being complementary to more standard map-based techniques rather than a replacement. Undoubtly,

there are analysis projects that are better suited to map-based techniques – typically those concerned with the science of local features in the maps, such as foregrounds. In addition, power spectrum estimation becomes cumbersome with purely harmonic methods if we have reason to question the foreground contamination of certain linear combinations of the spherical multipoles (e.g. those corresponding to features localised in the galactic plane).

The biggest hurdle facing a practical implementation of the harmonic ‘map’-making method at the resolution demanded by the upcoming satellite missions, is the need to solve the $O(l_{\text{max}}^2) \times O(l_{\text{max}}^2)$ linear system in equation (28). Unlike conventional map-making in the presence of noise correlations, the problem lies not in the inversion of the noise covariance matrix, since we are working in a representation where this matrix is already very sparse. The inversion of $\mathbf{A}^\dagger \mathbf{N}^{-1} \mathbf{A}$ can be avoided by adopting iterative techniques with a block-diagonal preconditioner. The main computational overhead arises from computing and storing the large, non-sparse matrix \mathbf{A} and applying it to the sky multipoles and its transpose to the Fourier ring data. These operations can be performed very efficiently using fast Fourier transform techniques for the rather idealised case of constant latitude scanning, with rings uniformly spaced in azimuth and no variations in instrument properties during the mission (Wandelt & Górski 2001). Unfortunately, it does not appear that such techniques can be easily extended to more realistic scan strategies. An assessment of some of these numerical problems, together with a number of potential solutions, will be given in Mortlock et al. (in preparation).

ACKNOWLEDGMENTS

This work benefited from useful discussions with several members of the Cambridge Planck Analysis Centre, in particular Martin Bucher, Rob Crittenden, and Neil Turok. ADC acknowledges a PPARC Postdoctoral Fellowship; DJM and MAJA are funded by PPARC.

REFERENCES

- Bond J. R., Jaffe A. H., Knox L., 1998, *Phys. Rev. D*, 57, 2117
 Bond J. R., Crittenden R. G., Jaffe A. H., Knox L., 1999, *Computing in Science and Engineering*, 1(2), 21
 Borrill J., 1999, in Maiani L., Melchiorri F., Vittorio N., eds, *3K cosmology: EC-TMR Conference*. American Institute of Physics, New York, p. 277
 Bouchet F. R., Prunet S., Sethi S. K., 1999, *MNRAS*, 302, 663
 Brink D. M., Satchler G. R., 1993, *Angular Momentum*, 3rd edn. Clarendon Press, Oxford
 Burigana C., Natoli P., Vittorio N., Mandolesi N., Bersanelli M., 2001, *A&A*, submitted (astro-ph/0012273)
 Challinor A., Fosalba P., Mortlock D., Ashdown M., Wandelt B., Górski K., 2000, *Phys. Rev. D*, 62, 123002
 Coble K. et al., 1999, *ApJ*, 519, L5
 de Bernardis P. et al., 2000, *Nature*, 404, 955
 Couchot F., Delabrouille J., Kaplan J., Revenu B., 1999, *A&AS*, 135, 579
 Delabrouille J., 1998, *A&AS*, 127, 555
 Delabrouille J., Górski K. M., Hivon E., 1998a, *MNRAS*, 298, 445
 Delabrouille J., Gispert R., Puget J.-L., 1998b, in Trần Thanh Vân T., Giraud-Héraud Y., Bouchet F., Damour T., Mellier Y. eds, *Fundamental Parameters in Cosmology*. Editions Frontieres, Paris, p. 263
 de Oliveira-Costa A., Kogut A., Devlin M. J., Netterfield C. B., Page L. A., Wollack E. J., 1997, *ApJ*, 482, L17
 Gaertner S. et al., 1997, *A&AS*, 126, 151
 Goldberg J. N., Macfarlane A. J., Newman E. T., Rohrlach F., Sudarshan E. C. G., 1967, *J. Math. Phys.*, 8, 2155
 Górski K. M., 1994, *ApJ*, 430, L85
 Hanany S. et al., 2000, *ApJ*, 545, L5
 Hobson M. P., Jones A. W., Lasenby A. N., Bouchet F. R., 1998, *MNRAS*, 300, 1
 Hobson M. P., Barreiro R. B., Toffolatti L., Lasenby A. N., Sanz J. L., Jones A. W., Bouchet F. R., 1999, *MNRAS*, 306, 232
 Janssen M. A. et al., 1996, preprint (astro-ph/9602009)
 Knox L., 1995, *Phys. Rev. D*, 52, 4307
 Kamionkowski M., Kosowsky A., Stebbins A., 1997, *Phys. Rev. D*, 44, 7368
 Lee A. T. et al., 2001, *ApJ*, submitted (astro-ph/0104459)
 Leitch E. M. et al., 2001, *ApJ*, submitted (astro-ph/0104488)
 Mortlock D. J., Challinor A. D., Hobson M. P., 2001, *MNRAS*, submitted (astro-ph/0008082)
 Natoli P., de Gasperis G., Gheller C., Vittorio N., 2001, *A&A*, 372, 346
 Netterfield C. B., Devlin M. J., Jarosik N., Page L., Wollack E. J., 1997, *ApJ*, 474, 47
 Netterfield C. B. et al., 2001, *ApJ*, submitted (astro-ph/0104460)
 Ng K., Liu G., 1999, *Int. J. Mod. Phys. D*, 8, 61
 Oh S. P., Spergel D. N., Hinshaw G., 1999, *ApJ*, 510, 551
 Padin S. et al., 2001, *ApJ*, 549, L1
 Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 1992, *Numerical Recipes: The Art of Scientific Computing*, 2nd edn. Cambridge Univ. Press, Cambridge
 Prunet S., Teyssier, R., Scully S. T., Bouchet F. R., Gispert R., 2001, *A&A*, 373, L13
 Revenu B., Kim A., Ansari R., Couchot F., Delabrouille J., Kaplan J., 2000, *A&AS*, 142, 499
 Scott P. F. et al., 1996, *ApJ*, 461, L1
 Smoot G. F. et al., 1992, *ApJ*, 396, L1
 Stolyarov V., Hobson M. P., Ashdown M. A. J., Lasenby A. N., 2001, *MNRAS*, submitted (astro-ph/0105432)
 Tanaka S. T. et al., 1996, *ApJ*, 468, L81
 Tegmark M., Efstathiou G., 1996, *MNRAS*, 281, 1297
 Tegmark M., 1997, *Phys. Rev. D*, 56, 4514
 Tegmark M., de Oliveira-Costa A., 2001, *Phys. Rev. D*, 64, 063001
 van Leeuwen F. et al., 2001, *MNRAS*, in press (astro-ph/0112276)
 Varshalovich D. A., Moskalev A. N., Khersonskii V. K., 1988, *Quantum Theory of Angular Momentum*. World Scientific, Singapore
 Vielva P., Barreiro R. B., Hobson M. P., Martinez-Gonzalez E., Lasenby A. N., Sanz J. L., Toffolatti L., 2001, *MNRAS*, submitted (astro-ph/0105387)
 Wandelt B. D., Hansen F. K., 2001, *Phys. Rev. D*, submitted (astro-ph/0106515)
 Wandelt B. D., Górski K. M., 2001, *Phys. Rev. D*, 63, 123002
 Wandelt B. D., Hivon E., Górski K. M., 2001, *Phys. Rev. D*, 64, 083003
 Wilson G. W. et al., 2000, *ApJ*, 532, 57
 Wu J. H. P. et al., 2001, *ApJS*, 132, 1

APPENDIX A: MAIN BEAM VARIATION ACROSS A SPECTRAL BAND

In this appendix we discuss the systematic error that arises from ignoring the variation of the beam multipoles with frequency across a spectral band when analysing microwave

data. This question does not appear to have received much attention so far in the CMB literature. For simplicity, we consider a single detector which is only sensitive to the total intensity. We take the beam profile to be an axisymmetric Gaussian, with frequency-dependent beam width $\sigma(\nu)$, so that for $\sigma(\nu) \ll 1$ we have

$$b_{(lm)}^I(\nu) = \delta_{m0} \sqrt{\frac{2l+1}{4\pi}} W_l(\nu), \quad (\text{A1})$$

where the scalar window function is

$$W_l(\nu) = \exp[-l(l+1)\sigma^2(\nu)/2]. \quad (\text{A2})$$

Assuming that the beam is diffraction limited, we expect the frequency dependence of the beam width to go like $\sigma(\nu) = (\nu_0/\nu)\sigma_0$, where $\sigma_0 \equiv \sigma(\nu_0)$ is the beam width at the central frequency of the band. For $l \gg 1$ this gives

$$b_{(lm)}^I(\nu) = \sqrt{\frac{\nu}{\nu_0}} b_{(lm)}^I(\nu_0/\nu, m), \quad (\text{A3})$$

which also holds more generally for arbitrary beams which scale inversely with frequency along longitudes.

If the CMB is the dominant physical component in the frequency band under consideration, the frequency spectrum of the brightness anisotropies in linear theory is

$$a_{(lm)}^I(\nu) = \tilde{a}_{(lm)}^I \bar{T} \left. \frac{\partial B(\nu, T)}{\partial T} \right|_{T=\bar{T}}, \quad (\text{A4})$$

for $l > 0$, where $B(\nu, T)$ is the Planck function and \bar{T} is the average CMB temperature over the sky. The $\tilde{a}_{(lm)}^I$, which describe the dimensionless anisotropy in the thermodynamic temperature of the sky, are independent of frequency for the CMB. We further restrict attention to white noise, assume uniform coverage of the spin axis pointing over the entire sky, and ignore any systematic variation in focal plane geometry. Then, if we attempt to solve for the $\tilde{a}_{(lm)}^I$, but ignore the variation of $W_l(\nu)$ with frequency, the maximum-likelihood solution returns a biased estimate $\hat{a}_{(lm)}^I$ which has the form

$$\hat{a}_{(lm)}^I = (1 + \delta_l) \tilde{a}_{(lm)}^I + n_{(lm)}, \quad (\text{A5})$$

where the fractional bias

$$\delta_l \equiv \frac{\int W_l(\nu) (\partial B / \partial T) v(\nu) d\nu}{W_l(\nu_0) \int (\partial B / \partial T) v(\nu) d\nu} - 1, \quad (\text{A6})$$

and the dimensionless random errors $n_{(lm)}$ have zero mean, and covariance

$$\langle n_{(lm)} n_{(l'm')}^* \rangle = \delta_{ll'} \delta_{mm'} w^{-1} W_l^{-2}(\nu_0). \quad (\text{A7})$$

Here, w is the dimensionless weight per solid angle (Knox 1995; see also Section 3). We can use the $\hat{a}_{(lm)}^I$ to estimate the CMB power spectrum C_l with the estimator

$$\hat{C}_l = \frac{1}{2l+1} \sum_{|m| \leq l} |\hat{a}_{(lm)}^I|^2 - w^{-1} W_l^{-2}(\nu_0), \quad (\text{A8})$$

which would be unbiased if δ_l were zero. The fractional systematic error in our estimate of C_l is therefore

$$(\Delta C_l / C_l)_{\text{syst}} = \delta_l (\delta_l + 2), \quad (\text{A9})$$

while there is a fractional random error

$$(\Delta C_l / C_l)_{\text{rand}} = \sqrt{\frac{2}{2l+1}} [1 + C_l^{-1} w^{-1} W_l^{-2}(\nu_0)]. \quad (\text{A10})$$




Figure A1. The systematic error in the recovered temperature power spectrum from neglecting the variation of the beam size across the spectral band for the 100 (solid lines), 143 (dashed lines), and 217 GHz (dash-dotted) *Planck* HFI channels. Also plotted is the random error in the C_l from the combined instrument (white) noise from all detectors at the specified frequency, and from cosmic variance. We have assumed that all beams are axisymmetric with Gaussian profiles, and that the beam size is diffraction limited.

The first term in brackets on the right-hand side of equation (A10) is cosmic variance. If instead we asked how well we could reconstruct the rotational-invariant $\sum_m |\tilde{a}_{(lm)}^I|^2 / (2l+1)$ for our given realisation of the sky, the cosmic variance term would not be present. The latter situation is the more relevant for mapping other astrophysical foreground components.

In Fig. A1 we plot the systematic and random errors on the CMB power spectrum for the 100, 143, and 217 GHz channels of the *Planck* HFI, using the predicted instrument specifications. The spectral filter $v(\nu)$ is assumed to be a top hat with fractional width $\Delta\nu/\nu_0 = 0.33$. It is clear that in the 100 GHz channel, the systematic error due to our neglect of the variation of beam width across the spectral band is non-negligible compared to the random error due to instrument noise over a broad range of l . For an axisymmetric beam the effect can be easily accounted for by including the frequency-dependent window function $W_l(\nu)$ in the integral over frequency, since the integral $\int a_{(lm)}^I(\nu) \tilde{b}_{d(lm')}^P(\nu) v(\nu) d\nu$ would still factor into a part depending on the sky with indices l and m only, as required for subsequent analysis. However, for a non-axisymmetric beam where the variation of $\tilde{b}_{d(lm')}^P(\nu)$ cannot be reduced to a frequency-dependent factor with only an l index, and a constant part with l and m' indices, we can no longer integrate the variation of beam with frequency and still preserve the factorisation of the sky. Note that if there is only an effective beam asymmetry, arising from skewing a symmetric beam with the temporal response of the instrument, the above comments do not apply since the frequency dependence of the effective beam will factor.

APPENDIX B: CONNECTING RING-SETS WITH MAPS

In this appendix we establish the result given as equation (66) in the main text, which relates the geometry of the ring-set to the observing time per solid angle on the sky.

The first step is to replace the single summation $\sum_{|n| \leq \min(l, l')}$ on the left-hand side of equation (66) with the double summation $\sum_{l'' \geq \max(|m|, |s|)} \delta_{l'l''} \sum_{|n| \leq \min(l, l'')}$, and to replace the labels l' by l'' in the argument of the summation. Denoting the left-hand side of equation (66) by ${}_s I_{mm'}^{ll'}$, we now have

$${}_s I_{mm'}^{ll'} = \sum_{l''n} [\delta_{l'l''} \delta_{mm'} \sqrt{(2l+1)(2l''+1)} \times d_{mn}^l(\theta) d_{mn}^{l''}(\theta) d_{ns}^l(\alpha) d_{ns}^{l''}(\alpha)], \quad (\text{B1})$$

where we have suppressed the limits on the summations. We now use the orthonormality of the spin-weighted harmonics,

$$\int -{}_s Y_{(l'm')}(\mathbf{e}_p) -{}_s Y_{(l''m)}^*(\mathbf{e}_p) d\Omega_p = \delta_{mm'} \delta_{l'l''}, \quad (\text{B2})$$

to replace the Kronecker deltas in equation (B1) with the integral of spin-weight $-s$ harmonics over the sphere. The term $-{}_s Y_{(l''m)}^*(\mathbf{e}_p)$ can be expressed in terms of D -matrices using equation (47):

$$-{}_s Y_{(l''m)}^*(\mathbf{e}_p) = (-1)^m \sqrt{\frac{2l''+1}{4\pi}} D_{sm}^{l''}(0, \theta_p, \phi_p), \quad (\text{B3})$$

where we have also used $d_{mm'}^l(\beta) = (-1)^{m-m'} d_{m'm}^l(\beta)$ (e.g. Brink & Satchler 1993), so that

$${}_s I_{mm'}^{ll'} = (-1)^m \sqrt{\frac{2l+1}{4\pi}} \int \sum_{l''n} \{e^{-im\phi_p} -{}_s Y_{(l'm')}(\mathbf{e}_p) \times [(2l''+1) d_{mn}^{l''}(\theta) d_{ns}^{l''}(\alpha) d_{sm}^{l''}(\theta_p)] \times d_{mn}^l(\theta) d_{ns}^l(\alpha)\} d\Omega_p. \quad (\text{B4})$$

If we now reverse the order of summation in equation (B4) using

$$\sum_{l'' \geq \max(|m|, |s|)} \sum_{|n| \leq \min(l, l'')} = \sum_{|n| \leq l} \sum_{l'' \geq \max(|m|, |n|, |s|)}, \quad (\text{B5})$$

the summation over l'' can be performed with the Ponzano-Regge sum rule (e.g. Varshalovich, Moskalev & Khersonskii 1988; equation 21, p. 89):

$$\begin{aligned} & \sum_{l'' \geq \max(|m|, |n|, |s|)} (2l''+1) d_{mn}^{l''}(\theta) d_{ns}^{l''}(\alpha) d_{sm}^{l''}(\theta_p) \\ &= \frac{2}{\pi} \frac{\Theta(B_p)}{\sqrt{B_p}} \cos(m\delta_1 + n\delta_2 + s\delta_3), \end{aligned} \quad (\text{B6})$$

where B_p is given by equation (61), and the angles δ_1 , δ_2 , and δ_3 are defined by the SO(3) composition

$$D(\delta_1, \theta, 0) D(\delta_2, \alpha, \delta_3) = D(0, -\theta_p, 0). \quad (\text{B7})$$

(The angles are given explicitly by Varshalovich et al. 1988, but it is straightforward to show that their definitions are equivalent to the implicit definition given here.) Finally, we can perform the remaining sum over n in equation (B4) by writing $\cos(m\delta_1 + n\delta_2 + s\delta_3)$ as the real part of $\exp[-i(m\delta_1 +$

$n\delta_2 + s\delta_3)]$, and combining the complex exponentials with the remaining d -functions to get

$$\begin{aligned} & \sum_{|n| \leq l} d_{mn}^l(\theta) d_{ns}^l(\alpha) \cos(m\delta_1 + n\delta_2 + s\delta_3) \\ &= \Re \sum_{|n| \leq l} D_{mn}^l(\delta_1, \theta, 0) D_{ns}^l(\delta_2, \alpha, \delta_3) \\ &= d_{sm}^l(\theta_p), \end{aligned} \quad (\text{B8})$$

where we have used equation (B7) in the last equality. The remaining terms in equation (B4) combine with $d_{sm}^l(\theta_p)$ to give $-{}_s Y_{(lm)}^*(\mathbf{e}_p)$ on using equation (B3). Our final result is

$${}_s I_{mm'}^{ll'} = \int \frac{2}{\pi} \frac{\Theta(B_p)}{\sqrt{B_p}} -{}_s Y_{(lm)}^*(\mathbf{e}_p) -{}_s Y_{(l'm')}(\mathbf{e}_p) d\Omega_p, \quad (\text{B9})$$

which establishes equation (66).

This paper has been produced using the Royal Astronomical Society/Blackwell Science L^AT_EX style file.